



HBAnyware® Utility

Version 4.0

User Manual

Copyright © 2008 Emulex. All rights reserved worldwide. No part of this document may be reproduced by any means or translated to any electronic medium without the prior written consent of Emulex.

Information furnished by Emulex is believed to be accurate and reliable. However, no responsibility is assumed by Emulex for its use; or for any infringements of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright or related rights of Emulex.

Emulex, AutoPilot Installer, BlockGuard, cLAN, FabricStream, FibreSpy, Giganet, HBAnyware, InSpeed, IntraLink, LightPulse, MultiPulse, SAN Insite, SBOD and Vixel are registered trademarks, and AutoPilot Manager, Critical Connectivity Solutions, EZPilot, SLI and VMPilot are trademarks of Emulex. All other brand or product names referenced herein are trademarks or registered trademarks of their respective companies or organizations.

Emulex provides this manual “as is” without any warranty of any kind, either expressed or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. Emulex may make improvements and changes to the product described in this manual at any time and without any notice. Emulex assumes no responsibility for its use, nor for any infringements of patents or other rights of third parties that may result. Periodic changes are made to information contained herein; although these changes will be incorporated into new editions of this manual, Emulex disclaims any undertaking to give notice of such changes.

Emulex, 3333 Susan Street
Costa Mesa, CA 92626

Introduction.....	1
Known Issues	3
Installing HBAnyware Components.....	5
Installing the HBAnyware Utility.....	5
In Windows	5
In Solaris LPFC, Solaris SFS, Linux and VMware ESX.....	5
Installing the HBAnyware Utility with Web Launch.....	8
Prerequisites	8
Procedure	9
Installing the HBAnyware CLI	10
Introduction	10
In Windows	10
In Linux	11
Installing the HBAnyware CLI on a Linux System With an Existing HBAnyware CLI Kit Installed	11
In VMware	12
Installing the HBAnyware CLI on a VMware System with an Existing HBAnyware CLI Kit Installed	12
Upgrading from CLI to Full-Featured HBAnyware	12
In Windows	12
In Linux	13
In VMware	13
Installing the HBAnyware Utility Security Configurator	13
Uninstalling the HBAnyware Security Configurator.....	14
Uninstalling HBAnyware Web Launch Only	14
Uninstalling the Utility Package	15
Changing Management Mode/Read-Only Mode.....	16
Using the HBAnyware Components.....	18
Starting the HBAnyware Utility.....	18
Starting the HBAnyware Utility with Web Launch.....	18
The HBAnyware Utility Window Element Definitions	19
The Menu Bar	20
The Toolbar	20
The Toolbar Buttons	20
The Discovery-Tree.....	21
Property Tabs.....	22
Status Bar	23
Customizing Tab Views	23
Discovering HBAs	24
Automatic Fibre Channel Discovery	24
Changing Adapter Port Names.....	25
Remote SAN Management Using TCP/IP Access Protocol.....	26
The Hosts File	26
Manually Editing the Hosts File	29
Copying the File	30
Configuring Discovery and TCP/IP Settings.....	30
Resetting Adapter Ports	31
Viewing Host Information.....	32
The Host Information Tab	32
The Driver Parameters Tab	33
Viewing Fabric Information	34
Viewing Virtual Port Information	35

Viewing Discovery Information	36
Viewing Adapter Information	37
Viewing Port Information	38
Viewing Port Statistics	40
Viewing Fabric Discovery Information	42
Viewing Transceiver Information	43
Viewing Vital Product Data (VPD).....	44
Viewing Maintenance Information.....	45
Viewing Target Information	47
Viewing LUN Information	48
Viewing Target Mapping (Windows, Solaris LPFC and Solaris SFS)	49
Viewing Target Mapping (Linux)	50
Creating and Deleting Virtual Ports	51
Creating Virtual Ports	51
Deleting Virtual Ports	53
Configuring the Driver	54
Setting Driver Parameters	54
Creating a Batch Mode Driver Parameters File	58
Assigning Batch Mode Parameters	60
Storport Miniport Driver Parameters	61
Server Performance (Windows)	67
Driver for Solaris LPFC – Configuration File Reference	69
Driver For Solaris SFS Driver Parameters	84
Solaris SFS and Solaris LPFC Driver Parameter Cross-Reference	90
Driver for Linux Parameter Tables	96
Version 8.2 LPFC Parameters.....	99
Driver for VMware ESX Configuration Parameters.....	104
Configuring Boot from SAN	106
Boot Types	106
Boot Device Parameters	107
Configuring Advanced Settings (Boot from SAN)	110
Using FC-SP DHCHAP Authentication (Windows, Linux 8.2, Solaris LPFC and Solaris SFS)	114
Changing Authentication Configuration	115
Changing Your Password	116
Viewing the Error and Event Log.....	116
Updating Firmware	116
Updating Adapter Firmware	116
Updating Adapter Firmware using Batch Mode	118
Downloading Converged Enhanced Ethernet Firmware (CEE)	120
Updating CEE Firmware using Batch Mode	124
Configuring CEE-Specific Parameters	128
Exporting SAN Information	129
Mapping and Masking	129
Automapping SCSI Devices (Windows)	129
Mapping and Masking Defaults (Windows)	130
Masking and Unmasking LUNs (Windows, Solaris LPFC and Solaris SFS)..	130
Using Automapping and Persistent Binding (Windows, Solaris LPFC and Solaris SFS).....	132
Diagnostics.....	135
Viewing Flash Contents, PCI Registers and Wakeup Information	136
Using Beaconing	138
Creating Diagnostic Dumps.....	138

Running Advanced Diagnostic Tests	140
Saving the Log File	143
Changing World Wide Name Configuration	144
HBAware Security.....	148
Introduction	148
Starting the HBAware Security Configurator	148
Running the Configurator for the First Time/Creating the ACG.....	149
Designating a Master Security Client.....	150
Access Control Groups.....	151
Introduction	151
Access Control Group Tab on the MSC.....	151
Access Control Group Tab on a Non-MSC	152
ACG Icons.....	152
Adding a Server to the ACG	153
Deleting a Server from the ACG.....	153
Removing Security from all Servers in the ACG.....	153
Generating New Security Keys	153
Restoring the ACG to Its Last Saved Configuration	154
Accessing a Switch	154
Access Sub-Groups.....	154
Introduction	154
ASG Icons	155
Creating an ASG	155
Reserved Indices - Examples.....	157
Adding a Server to an ASG	157
Deleting an ASG	157
Restoring an ASG to Its Last Saved Configuration.....	157
Editing an ASG	158
About Offline ASGs	159
Backup Masters.....	160
Introduction	160
Backup Master Eligible Systems	161
Backup Master Tab and Controls	161
Creating a Backup Master.....	162
Reassigning a Backup Master as the New MSC from the Old MSC.....	162
Reassigning a Backup Master as the New MSC from the Backup Master	163
Using the HBAware Utility Command-Line Interface.....	164
Using the CLI Client	164
Syntax Rules	164
The CLI Client Command Reference.....	165
Read-Only Mode	165
Help Commands	165
Attributes Commands.....	166
Authentication Commands	167
Boot Commands	169
CEE Commands	170
Diagnostic Commands	171
Driver Parameter Commands.....	174
Dump Commands	176
LUN Masking Commands	177
Miscellaneous Commands	179

Persistent Binding Commands	180
TCP/IP Management Host File Commands.....	182
VPort Commands.....	183
WWN Management Commands	184
Troubleshooting	186
General Situations	186
Emulex Driver for Windows and the HBAnyware Utility Situations	188
Emulex Driver for Solaris LPFC and the HBAnyware Utility Situations	188
Emulex Driver for Linux and the HBAnyware Utility Situations	189
VPorts and the HBAnyware Utility Situations.....	194
Security Configurator Situations - Access Control Groups (ACG)	194
Security Configuration Situations - Access Sub-Groups (ASG)	195
HBAnyware Security Configurator Situations - Backup Masters	196
Error Message Situations	197
Master Security Client Situations.....	198

Introduction

The HBAAnyware® utility is a powerful, centralized adapter management suite, providing discovery, reporting and management of local and remote adapters from a single console anywhere in the SAN and across platforms. Both a graphical user interface (GUI) and command line interface (CLI) are provided. This remote configuration capability can be provided by either Fibre Channel (FC) access via host systems on the same FC Storage Area Network (SAN) or by Transmission Control Protocol/Internet Protocol (TCP/IP) access from IP addresses of remote machines.

This manual supports the following versions of the HBAAnyware utility:

- Windows
- Solaris LPFC
- Solaris SFS ('emlxs' is the module name for the Emulex driver for Solaris SFS)
- Linux
- VMware ESX Server

Use the HBAAnyware utility to do any of the following (refer to Table 1 to determine if a specific feature or task is supported by your operating system):

- Discover local and remote hosts, adapters, targets and Logical Unit Numbers (LUNs)
- Enable local and FC discovery of Emulex and OEM branded Emulex adapters
- Change an adapter's World Wide Port Name (WWPN) or World Wide Node Name (WWNN) (new in version 4.0)
- Reset adapters
- Set up persistent binding
- Set adapter driver parameters simultaneously to multiple adapters using Batch Update
- Set global driver parameters for adapters
- Update firmware and FC boot code (x86 BootBIOS, OpenBoot or EFIBoot) on a single adapter or multiple adapters using Batch Update
- Enable or disable the adapter BIOS (x86 BootBIOS, FCode or EFIBoot)
- Run diagnostic tests on adapters
- Manage local, FC remote and TCP/IP-accessed adapters
- Locate adapters using beaconing
- Mask and unmask LUNs
- Perform authentication using the Fibre Channel Security Protocol Diffie-Hellman Challenge Handshake Authentication Protocol (FC-SP DHCHAP)
- Create and delete virtual ports (N_Port_ID virtualization [NPIV] must be enabled.) (new in version 4.0)
- Run in read-only mode (new in version 4.0)
- Configure boot from SAN (new in version 4.0)
- Modify an IP port number (new in version 4.0)
- View vital product data (VPD) for the selected adapter port (new in version 4.0)
- View transceiver information for the selected adapter port (new in version 4.0)
- Create reports about discovered SAN elements (new in version 4.0)

Table 1: The HBAnyware Utility Features and Tasks Cross-Reference

Feature/Task	Windows	Solaris LPFC	Solaris SFS	Linux	VMware ESX Server
HBAnyware Graphical User Interface (GUI)	X	X	X	X	X*
HBAnyware Command Line Interface (CLI)	X	X	X	X	X
HBAnyware with Web Launch utility	X	X	X	X	
HBAnyware Security Configurator	X	X	X	X	
Discover local hosts, adapters, targets and LUNs	X	X	X	X	X*
Discover remote hosts, adapters, targets and LUNs	X	X	X	X	X*
Enable local discovery of Emulex and OEM branded Emulex adapters	X	X	X	X	X*
Enable FC discovery of Emulex and OEM branded Emulex adapters	X	X	X	X	X*
Change an adapter's WWPN or WWNN	X	X	X	X	X*
Reset adapters	X	X	X	X	X*
Set up persistent binding	X	X	X		
Set adapter driver parameters simultaneously to multiple adapters using Batch Update	X	X	X	X	
Set global driver parameters to adapters	X	X	X	X	X**
Boot from SAN functionality	X	X	X	X	X
Update firmware and FC boot code on a single adapter or multiple adapters using Batch Update	X	X	X	X	X*
Enable or disable the x86 BootBIOS	X	X	X	X	X*
Run diagnostic tests on adapters	X	X	X	X	
Manage local adapters	X	X	X	X	X*
Manage FC remote and TCP/IP accessed adapters	X	X	X	X	X*
Locate adapters using beaconing	X	X	X	X	X
Mask and unmask LUNS	X	X	X		
Perform authentication using FC-SP DHCHAP	X	X	X	X	

* Supported only by hbacmd on for the VMware release of the HBAnyware utility 4.0. Remote management clients can perform these functions on ESX Server HBAs using the HBAnyware GUI.

** Temporary (not persistent) driver parameters are supported only by hbacmd on VMware ESX Server.

Known Issues

The following issues have been reported at the time of publication. These issues may not have been verified or confirmed and can apply to another product, such as hardware.

- Emulex provides support for LightPulse® adapters that are reprogrammed with WWPNS outside the typical Emulex range, such as Hewlett-Packard's Virtual Connect for FC on the BladeSystem c-Class platform. In these environments, the HBAnyware utility must be deployed across all servers on the SAN, and on any other management console used for TCP/IP access management so that all adapters appear in the discovery-tree.
- Multiport adapter models (LP9802 and later) are represented in the discovery-tree as a single adapter with separate port icons for each port. Older multiport models (for example LP8000DC, LP9402DC or LP9002DC) are represented by separate adapter icons.
- If there are multiple versions of the Java Runtime Environment (JRE) installed on your Internet Explorer client, you may see the following text in the browser's main display window when you attempt to launch the HBAnyware utility via the browser:

```
Emulex Corporation HBAnyware Demo of HBAnyware WebStart web n.n.n.n.....
```

If you have verified that the HBAnyware Web Launch Service package is installed and is running on the target server, try one of these two workarounds:

- Exit the browser, then restart it. The HBAnyware utility launches successfully.
- Uninstall all non-essential versions of the JRE. HBAnyware Web Launch Service requires only a single version of the JRE be installed on the Windows browser client.
- The JRE version required for the HBAnyware GUI is no longer supported on Linux Itanium (IA64) systems. The HBAnyware CLI is still supported on these systems.
- On VMware ESX Servers there is an issue with the 'discovery-threads' parameter. The range for this parameter should be 30-64 (decimal). The default is erroneously set to 1. Set this driver parameter to a valid value. If you do not, when you change any driver parameter value, the following error message may be displayed:

```
Driver Parameter 'discovery-threads' is not within the allowed range.
```

- For Solaris SFS, COMSTAR Beta Support:
 - NPIV is not yet supported on target mode ports.
 - DHCHAP is not yet supported on target mode ports.
- Internal and External Loopback diagnostic tests are not available for LP21000 and LP21002 adapters. Centralized management of adapters across VMware ESX servers must be limited to at most one HBAnyware client. This limitation applies to the HBAnyware utility remote connectivity for both FC (in-band) and TCP/IP (out-of-band) management. The remote HBAnyware client configuration must be modified to disable automatic periodic polling of remote servers.

To disable automatic periodic polling of remote servers:


1. From the **Discovery** menu, click **Modify Settings**. The HBA Discovery Properties dialog box appears.
2. Define the discovery properties as follows:
 - Discovery Server Startup area: This setting defaults to **When this utility starts**. Keep this default or if necessary select it at system boot.
 - Refresh Rate, Fibre Channel area: Select **Manual Refresh** (requires clicking **Discovery Refresh**).

- Refresh Rate, TCP/IP area: Select **Manual Refresh** (requires clicking **Discovery Refresh**).
 - Expire Undiscovered HBA area: Select **Never Remove**.
3. Click **OK**.
- LP21000 and LP21002 adapters support a link speed of 4 Gb/s only.
 - If you connect an LP21000 or LP21002 adapter to a Nuova switch, you can only select point-to-point topology.
 - Emulex recommends that remote management of other Windows, Linux and Solaris servers take place in an environment separate from management of ESX servers. For example, other servers should exist in a separate zone so that the HBAnyware client that manages them does not automatically discover ESX servers.
 - For Solaris SFS, the enable-npiv driver parameter is disabled by default and the following message is displayed on the Virtual Ports tab:

The fabric switch does not support virtual ports.

This message is displayed regardless of the switch's ability to support NPIV. Refer to your switch documentation and/or manufacturer to verify that NPIV is supported.

To enable this parameter using the HBAnyware utility:

1. Select your host from the discovery-tree.
 2. Select the host **Driver Parameters** tab.
 3. Highlight the **enable-npiv** parameter and select **Enabled**.
 4. From the discovery-tree, select a port that is connected to the switch.
 5. Reset the adapter: from the Port menu, click **Reset HBA Port** or click .
 6. When prompted, click **Yes** to continue.
- Switch support for NPIV - When the enable-npiv driver parameter is disabled, the Virtual Ports tab may erroneously report the following message:

The fabric switch does not support virtual ports.

To determine if the switch supports NPIV, either:

- Enable the enable-npiv driver parameter using the HBAnyware utility. (see “Enabling the enable-NPIV Parameter” below.)

Or:

- Refer to your switch documentation or manufacturer to determine if the switch supports NPIV.

Enabling the enable-NPIV Parameter

1. Select your host from the discovery tree.
2. Select the **Host Driver Parameters** tab.
3. Highlight **enable-npiv**.
4. Select **Enabled**.
5. From the discovery-tree, select a port that is connected to the switch.
6. **Reset HBA Port** to reset the HBA.
7. When prompted, press **Yes** to continue.

When the enable-npiv parameter is enabled, the following message is displayed only if the switch actually does not support NPIV:

The fabric switch does not support virtual ports.

- The global link-speed setting for initializing the FC connection cannot be changed from the default value "Auto-Detect" to 1, 2, 4, or 8 Gb/s.
- On Solaris LPFC, it is possible to mask and unmask LUNs, even if the HBAnyware utility is configured for read-only mode.
- While adding an out-of-band host in HBAnyware, if the utility does not recognize a valid hostname, verify that the hostname and corresponding IP address have properly been added to the /etc/hosts file. The HBAnyware GUI and HBAnyware CLI resolve hostnames by searching the /etc/hosts file.
- The global default driver parameters for converged network adapters (LP 21000 series CNAs) cannot be changed.

Installing HBAnyware Components

Installing the HBAnyware Utility

In Windows

The AutoPilot Installer® software streamlines the Emulex driver and HBAnyware utility installation. Refer to the Quick Installation Manual for more information. This manual is available on the Emulex Web site for your driver version.

In Solaris LPFC, Solaris SFS, Linux and VMware ESX

The following must be installed before you can install the utilities:

- The appropriate driver for your operating system:
 - Solaris LPFC driver version 6.20i or later.
 - Solaris SFS driver version 2.21 or later
 - Linux driver version 8.0.16.34 or later.
 - Linux driver version 8.2.0.25 or later
 - Emulex Driver for VMware ESX, version 7.4 or later. Refer to the Emulex Driver for VMware ESX User Manual for specific information on driver support in ESX Releases.
- For Solaris LPFC and Solaris SFS, JRE 5.0; HBAnyware utilities do not run under earlier versions of the JRE. The JRE and instructions for installation are available at <http://java.sun.com/downloads/index.html>.

Caution: The utilities require Java runtime binaries and libraries. Their paths must be included at the beginning of the PATH environment variable to avoid conflicts with earlier versions of Java that can still be installed on the system. For example, if the Java runtime binaries are in /usr/java/bin, then include this path in the PATH environment variable. For example: (bash> export PATH="/usr/java/bin:\$PATH")

- For Solaris SFS, the Emulex Fibre Channel Adapter (FCA) utilities; Refer to the FCA Utilities User Manual for instructions on unpacking and installing the FCA utilities.
- In Linux, previous versions of the application helper module must be uninstalled. You must run the uninstall script that shipped with the version of the application helper module you want to remove.

To install the HBAnyware utilities in Solaris LPFC and Solaris SFS:

1. Uncompress and untar the EmlxApps file included in the driver package. For Solaris SFS, proceed to step 3.
2. For Solaris LPFC, run the unpack script to obtain the correct package version. Type:
`./unpack_apps`
3. Unzip the file. Type:
`gunzip HBAnyware-<version>-<platform>.tar.gz`
4. Untar the file. Type:
`tar -xvf HBAnyware-<version>-<platform>.tar`
5. Run the pkgadd utility. Type:
`pkgadd -d .`
6. When prompted by pkgadd, choose to install the HBAnyware utilities.
7. When prompted by pkgadd, answer the HBAnyware installation option questions.

To install the HBAnyware utilities in Linux:

Note: The HBAnyware utility GUI and Security Configurator (SSC) GUI applications are not supported on Linux for the IA64 platform.

Note: For Linux 8.0: This procedure also installs the application helper module on your system. The application helper module allows HBAnyware to communicate with the Emulex driver for Linux. The 'elxlpfc' init script is also installed and configured to start and stop the 'lpfcdfc' driver during system startup and shutdown.

1. Log on as 'root'.
2. Download the utilities from the Emulex web site or copy them to the system from the installation CD.
3. Copy the installation and uninstallation scripts to a known location, for easy access by other users.
4. Copy the ElxLinuxApps-<AppsRev><DriverRev>.tar file to a directory on the install machine.
5. Change (use cd command) to the directory to which you copied the tar file.
6. Untar the file. Type:
`tar -xvf ElxLinuxApps-<AppsRev><DriverRev>.tar`
7. Uninstall any previously installed versions. Type:
`./uninstall`
8. Run the install script. Type:
`./install`
9. Enter the type of management you want to use:
 - 1 Local Mode : HBA's on this Platform can be managed by HBAnyware clients on this Platform Only.
 - 2 Managed Mode: HBA's on this Platform can be managed by local or remote HBAnyware clients.
 - 3 Remote Mode : Same as '2' plus HBAnyware clients on this Platform can manage local and remote HBA's.

10. If you answered <2> or <3> in step 8, you are asked if you want the HBAnyware utility to operate in read-only mode. Read-only mode prevents users from performing certain operations such as resetting HBAs, updating an adapter's firmware and changing adapter driver properties and bindings. Enter <y> 'for yes to allow the user to perform these operations, enter <n> for no if read-only mode is desired.
11. You are prompted as to whether or not to allow users to change the management mode after installation. Enter 'y' for yes, or 'n' for no.

You can also install the applications kit on an upgraded kernel. The lpfc driver must be part of the target kernel distribution and the utilities package must have been installed on the current kernel.

To install the applications kit on an upgraded kernel:

1. Boot to the new kernel.
2. Log on as 'root'.
3. Change (use the cd command) to the directory containing the unpacked Applications Kit.
4. Run the install upgrade kernel script. Type:
`./install upgradkernel`

To install the HBAware Agent in VMware ESX Server:

The LPFC driver must be loaded before you can install the HBAware Agent.

1. Log in as 'root'.
2. Copy the `elxvmwarecorekit-<AppsRev>.rpm` file to a directory on the install machine.
3. CD to the directory to which you copied the rpm file.
4. Install the rpm. Type:

```
rpm -ivh elxvmwarecorekit-<AppsRev>.rpm
```

The rpm contents are installed in `/usr/sbin/hbanyware`. The `hbacmd` utility is also located in this directory.

Installing the HBAware Utility with Web Launch

Prerequisites

In addition to the driver and HBAware utilities, the following prerequisites must be met before you install the Web Launch feature:

Note: The HBAware utility with Web Launch is not supported on VMWare ESX Server.

- In Windows:
 - Microsoft Internet Information Services (IIS) Server must be installed. See the Microsoft Web site for information on downloads and installation.
 - JRE must be installed. See the www.java.com Web site for information on downloads and installation.
 - The Windows Firewall feature may be enabled by default. If it is, you must add and enable three exceptions: HTTP port, `java.exe` and `rmiregistry.exe` (both included with the JRE).

Note: Allowing programs and/or ports through the firewall may increase the security risks. Use at your own discretion.

To enable the HTTP port:

1. Click **Add Port...** The Add a Port dialog is displayed.
2. On the Add a Port dialog, type `HTTP` as the Name and `80` as the Port Number.
3. Leave the radio button on **TCP** and click **OK**.

To enable the `java.exe` program:

1. Click **Add Program...** The Add a Program dialog is displayed.
2. Click **Browse...**
3. Specify `java.exe` located in the bin directory of the JRE installation path. Example:
`C:\Program Files\Java\jre1.5.0_06\bin\java.exe`.
4. Click **OK**.

To enable the `rmiregistry.exe` program:

1. Click **Add Program...** The Add a Program dialog is displayed.
 2. Click **Browse...** and specify `rmiregistry.exe` located in the bin directory of the JRE installation path. Example:
`C:\Program Files\Java\jre1.5.0_06\bin\rmiregistry.exe.`
 3. Click **OK**.
 4. Click **OK** to apply the new firewall settings.
- In Solaris LPFC, Solaris SFS and Linux:
 - Apache must be installed and running on the server that is hosting the Web Launch Service software.
 - The Java Web Start application must be installed and running on the browser host.

The system on which you are installing the Web Launch Service package (the server) requires:

- An HTTP server configured to handle the JNLP MIME file type. The following MIME file type/ file extension must be added to your server configuration:

```
MIME type: application/x-java-jnlp-file
File Extension: jnlp
```

- The HTTP server must be running.

The system on which you are running the browser (the client) requires:

- JRE 5.0 or later must be installed. The HBAware-installed JRE must match the HBAware code base. Specific requirements:
- Sun 32-bit JRE 5.0 or later for Intel based systems (x86 and IA64)
- Sun 32-bit JRE 5.0 or later x86-64
- 64-bit JRE 5.0 or later for RH4 and SL9 (ppc64)
- 32-bit JRE 5.0 or later for RH5 and SL10 (ppc64)

Refer to the appropriate vendor documentation for detailed instructions about configuring MIME types, configuring and starting the HTTP server and installing the JRE.

Procedure

To install HBAware with WebLaunch:

In Windows (Windows Server 2003, Windows Vista and Windows Server 2008):

Click **Programs>Emulex >HBAware WebLaunch Install**. Web Launch installation begins.

In Solaris LPFC, Solaris SFS and Linux:

1. Log on as 'root'.
2. Navigate to the HBAware directory.
 - Solaris LPFC and Solaris SFS:
`cd /opt/HBAware`
 - Linux:
`cd /usr/sbin/hbanyware`
3. Run the install script. Type:
`./wsinstall`
4. When prompted, enter the Web server's document root directory. For example:
`/srv/www/htdocs`

5. You are provided with the IP address of the host and asked if that is the IP address that the Web server uses. Answer Y or N as appropriate. If you answer N, you are prompted for the IP address you want to use.
6. You are asked if your Web server is listening on the normal default HTTP port (80). Answer <y> or <n> as appropriate. If you answer <n>, you are prompted for the port you want to use.

Once you have entered the necessary information, you are notified when the installation of the HBAnyware Web Launch package is complete. The Web Launch configuration files are created and Web Launch Service automatically starts.

7. To verify the installation, locate another client, open a Web browser window and enter this URL according to this format:

`http://IP_ADDR:PORT_NUM/hbanyware.jnlp`

where *IP_ADDR* is the IP address of host on which you installed the HBAnyware Web Launch service, and *PORT_NUM* is the TCP port number of the listening hosts' Web server. The standard HBAnyware user interface is displayed.

Note: It is not necessary to enter a port number if the standard HTTP port was chosen during configuration.

Installing the HBAnyware CLI

Introduction

The HBAnyware CLI is a separate application with core driver kits that do not include the HBAnyware GUI. The HBAnyware CLI console application name is `hbacmd` and can be installed on Windows, Linux and VMware. A single operation is performed by entering 'hbacmd' at the command line. For syntax information and details on using the HBAnyware CLI, see "Using the CLI Client" on page 164.

Platforms that are supported with the HBAnyware CLI are detailed in Table 2.

Table 2: HBAnyware Command Line Interface Supported Platforms

Driver	Architecture	Operating System
Storport Miniport Driver	Intel x86 and x64	Windows Server 2003
LPFC 7.4.x Driver	Intel x86, EM64T and AMD64	VMware ESX Server 3.5
Open Source Driver for RHEL4 and SUSE Linux Enterprise Server (SLES) 9	Intel x86, EM64T, AMD64, PPC64 and IA64	RHEL 4 Updates 3, 4 and 5 and later SUSE Linux Enterprise Server 9 SP 2 or later
Open Source Driver for RHEL 5 and SLES 10	Intel x86, EM64T, AMD64, PPC64 and IA64	RHEL 5, RHEL 5.1 and RHEL 5.2 SUSE Linux Enterprise Server

In Windows

To install the HBAnyware CLI, run an installation .exe file for a core Windows driver kit that does not include the HBAnyware GUI:

- `storportminiportcorekit_[version].exe`
[version] represents the complete version. For example, `storportminiportcorekit_1-30a9-1d`

In Linux

Before installing the core kit, the 8.0 or 8.2 driver must be installed. For 8.0 systems (RHEL4, SLES9), the `lpfcdfc` IOCTL module must also be installed. To install the HBAnyware CLI on a new system, install the specific driver RPM for your Linux version.

Enter the following command all in one line:

```
# rpm -i elxlinuxcorekit-[version].rpm
```

Installing the HBAnyware CLI on a Linux System With an Existing HBAnyware CLI Kit Installed

Follow these steps to install the HBAnyware CLI on a Linux system with an existing HBAnyware CLI kit installed:

1. Uninstall the Linux core kit. Type:

```
rpm -e elxlinuxcorekit-[version]
```

Note: If the uninstallation script does not work, you have an older HBAnyware kit. In this case, follow the procedure for **Uninstalling Older HBAnyware Kits on VMware** in this topic.

2. Install the specific RPM for your driver for Linux version. Enter this command (all in one line):

```
# rpm -i elxlinuxcorekit-[version].rpm
```

Uninstalling Older HBAnyware Kits on Linux

1. Locate and download the full application tar file for the appropriate Linux version.
2. Untar the tar file and run the installation script to install the application.

If the HBAnyware Security Configurator is installed, it must be uninstalled before uninstalling the HBAnyware utility. You must run the uninstall script that shipped with the version of HBAnyware Security Configurator that you want to remove. Proceed to step 3. If the Security Configurator is not installed, proceed to step 4.

3. If the HBAnyware Security Configurator is installed, follow these steps:
 - a. Log on as 'root'.
 - b. Change (use `cd` command) to the directory to which you copied the tar file during installation.
 - c. Run the uninstall script with the `ssc` parameter specified. Type:

```
./uninstall ssc
```
4. Uninstall the HBAnyware utility, `lputil` and the Application Helper Module:
 - a. Log on as 'root'.
 - b. Change (use `cd` command) to the directory to which you copied the tar file during installation.
 - c. Uninstall any previously installed versions. Type:

```
./uninstall
```
 - d. Install the specific RPM for your driver for Linux version. Enter the following command all in one line.

```
# rpm -i elxlinuxcorekit-[version].rpm
```

In VMware

To install the HBAnyware CLI on a new system, install the specific RPM for the driver for your VMware version.

Prerequisites

- The lpfc driver must be loaded.

Procedures

To install the HBAnyware CLI:

1. Log in as 'root'.
2. Copy the `elxvmwarecorekit-<kit version>.rpm` file to a directory on the install machine.
3. CD to the directory to which you copied the rpm file.
4. Install the rpm. Type:

```
rpm -i elxvmwarecorekit-<kit version>.rpm
```

The rpm contents are installed in `/usr/sbin/hbanyware`. The `hbacmd` utility is also located in this directory.

Installing the HBAnyware CLI on a VMware System with an Existing HBAnyware CLI Kit Installed

Follow these steps to install the HBAnyware CLI on a VMware system with an existing HBAnyware CLI kit Installed:

1. Install the RPM by entering the following command all on one line:

```
# rpm -U elxvmwarecorekit-[kit version].rpm
```

Uninstalling Older HBAnyware Kits on VMware

1. Log in as 'root'.
2. Type: `rpm -qa | grep elx` and locate the following rpm file:

```
elxvmwarecorekit-<kit version>
```

The rpm contents are installed in `/usr/sbin/hbanyware`. The `hbacmd` utility is also located in this directory.

3. Type:

```
rpm -e elxvmwarecorekit-<kit version>
```

Upgrading from CLI to Full-Featured HBAnyware

In Windows

To upgrade from the HBAnyware CLI to the full-featured HBAnyware utility:

1. From the desktop, run one of the .exe files that contain the full application kit:

- `storportminiportkit_[version].exe`
- `scsiportminiportkit_[version].exe`

[version] represents the complete version. For example, `storportminiportkit_1-30a9-1d`.

Running this executable file removes the HBAware CLI and installs a full-featured version of the HBAware utility that includes the CLI and the GUI.

In Linux

To upgrade from the HBAware CLI to the full-featured HBAware utility:

1. Uninstall the core kit, using `rpm -e elxlinuxcorekit-[version]`.
2. Install the HBAware kit, using the install script within the tar file.

In VMware

The full-featured HBAware kit is not supported on VMware ESX Server.

Installing the HBAware Utility Security Configurator

The Emulex driver and the HBAware utilities must be installed before you can install the HBAware Security Configurator.

Note: The HBAware utility Security Configurator is not supported on VMWare ESX Server.

To install the HBAware utility Security Configurator:

In Windows:

1. Locate the SSCsetup.exe file. The default path for this file is:
`C:\Program Files\Emulex\Util\HBAware`
2. Double-click the **SSCsetup.exe** file. A welcome window appears.
3. Click **Next**. The Setup Status window is displayed. After setup completes, the Emulex HBAware Security Setup Completed window appears.
4. Click **Finish**.

In Solaris LPFC and Solaris SFS:

1. Copy the <HBAwareSSC_version>.tar.gz file to a directory on the install machine.
2. cd to the directory to which you copied the .gz file.
3. Untar the file. Type: `gzcat <HBAwareSSC_version>.tar.gz | tar xvf-`
4. At the shell prompt, type:
`pkgadd -d `pwd``
5. When prompted by pkadd, choose to install HBAwareSSC.
6. When prompted by pkadd, answer the HBAware installation option questions.

In Linux:

1. Log on as 'root'.
2. Change (use the cd command) to the directory to which you copied the tar file. (See "Installing the Utilities and the application helper module" on page 7 step 2 for reference.)
3. Run the install script with the ssc parameter specified. Type:
`./install ssc`

Uninstalling the HBAware Security Configurator

To uninstall the HBAware Security Configurator:

In Windows:

1. Select **Start>Settings>Control Panel**. The Control Panel appears.
2. Click **Add/Remove Programs**. The Add or Remove Programs window appears.
3. Select **Emulex HBAware Security Configurator>Change/Remove**.
4. Click **Next**. The Security Configurator is removed from the system.
5. Click **Finish**. Uninstallation is complete.

In Solaris LPFC and Solaris SFS:

1. Log on as 'root'.

Note: If the HBAware Security Configurator is installed, it must be uninstalled before uninstalling the HBAware and driver utilities.

2. Type:

```
pkgrm HBAwareSSC
```

In Linux:

Note: You must run the uninstall script that shipped with the version of HBAware Security Configurator you want to remove. If the uninstall script resides in the usr/src directory, be sure to copy it to a temporary directory before you run it.

1. Log on as 'root'.
2. Change (use the cd command) to the directory to which you copied the tar file during installation.
3. Run the uninstall script with the ssc parameter specified. Type:

```
./uninstall ssc
```

Uninstalling HBAware Web Launch Only

To uninstall HBAware Web Launch, but leave the HBAware utility installed:

In Windows:

1. Select **Start> Programs>Emulex>HBAware WebLaunch Uninstall**. The following screen appears:

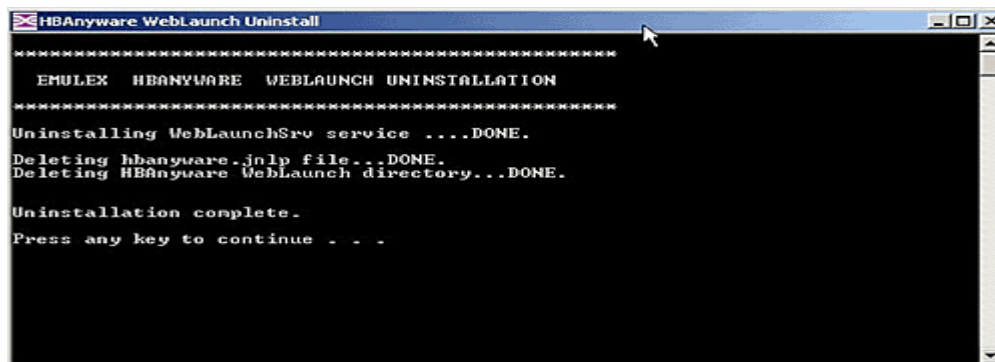


Figure 1: HBAware Web Launch, Uninstall screen

2. HBAware Web Launch is removed. Press any key to continue.

In Solaris LPFC, Solaris SFS and Linux:

1. Log on as 'root'.
2. Execute the uninstallation script.
 - Solaris LPFC and Solaris SFS:
`/opt/HBanyware/wsuninstall`
 - Linux:
`/usr/sbin/hbanyware/wsuninstall`

This script stops the HBAnyware Web Launch Service daemons (if they are running) and removes all Web Launch related files from the host.

Uninstalling the Utility Package

Note: If you installed HBAnyware with Web Launch, you must uninstall it before uninstalling the HBAnyware utility.

To uninstall the HBAnyware utility and HBAnyware Web Launch:

In Windows:

1. Select **Start>Settings>Control Panel**. The Add/Remove Programs window appears. Select the **Install/Uninstall** tab.
2. Select **Emulex HBAnyware** and click **Remove**. Click **Yes**. The utilities are removed from the system.
3. Select Emulex Common SAN Management and click **Remove**. Click **Yes**. The Emulex Common SAN Management components are removed from the system.
4. Click **Finish**. Uninstallation is complete.

In Solaris LPFC and Solaris SFS:

1. Log on as 'root'.
2. Type:
`pkgrm HBAnyware`

In Linux (also uninstalls the application helper module):

1. Log on as 'root'.
2. Change (use the `cd` command) to the directory to which you copied the tar file during installation.

Note: If you cannot find the original uninstall script, uninstall the HBAnyware utility and HBAnyware Web Launch by running:

`/usr/sbin/hbanyware/scripts/uninstall_hbanyware`

3. Uninstall any previously installed versions. Type:
`./uninstall`

In VMware ESX Server (uninstalls the HBAnyware Agent):

1. Log in as 'root'.
2. Type:
`rpm -qa | grep elx`

3. Locate the `elxvmwarecorekit-<AppsRev>.rpm` file. The `.rpm` contents are installed in `/usr/sbin/hbanyware`. The `hbacmd` utility is also located in this directory.
4. Type:


```
rpm -e elxvmwarecorekit-<kit version>
```

Changing Management Mode/Read-Only Mode

During installation, you selected both a management and a read-only mode. If you also chose to enable modification of these settings after installation, then you can choose three types of host/adapter management:

- **Strictly Local Management** - This setting only allows management of adapters on this host. Management of adapters on this host from other hosts is not allowed.
- **Local Management Plus** - This setting only allows management of adapters on this host, but management of adapters on this host from another host is possible.
- **Full Management** - This setting enables you to manage adapters on this host and other hosts that allow it.

If Management Mode was enabled during installation, you can also set read-only mode.

- **Read-only mode** - This setting prevents performance of certain operations such as resetting adapters, updating the adapter or CEE firmware image and changing adapter driver properties and bindings. Dialog box buttons and menus that pertain to these tasks are completely hidden or inactive.

To change management/read-only mode:

Note: You must restart the HBAnyware utility to see the new management mode.

In Windows:

1. From the **File** menu, select **Management Mode**. The Management Mode dialog box appears.

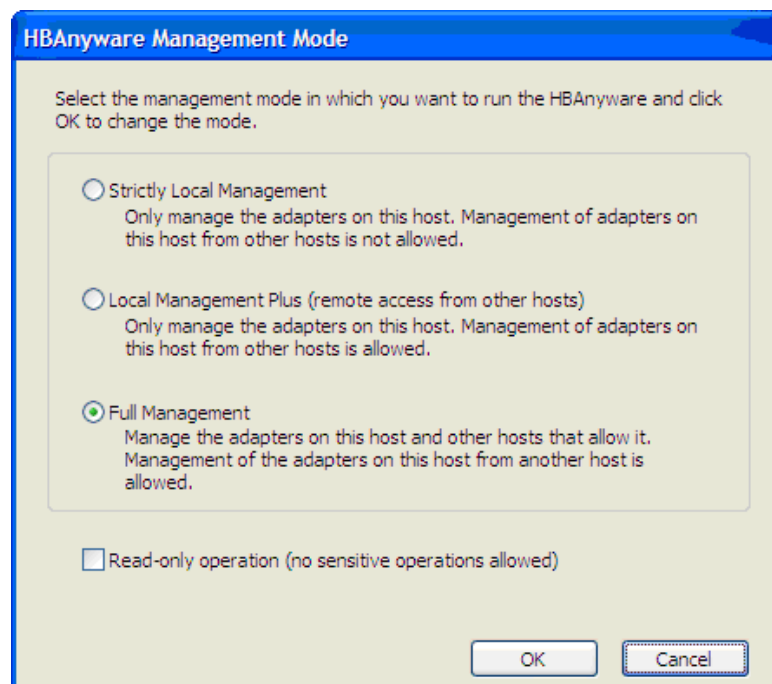


Figure 2: Management Mode dialog box

2. Choose the management type and read-only mode you want.
3. Click **OK**.

In Solaris LPFC and Solaris SFS:

1. Run the following script:
`cd /opt/HBAnyware/set_operating_mode`
2. Choose the management type and read-only mode you want.

In Linux:

1. Run the following script:
`cd /usr/sbin/hbanyware/set_operating_mode`
2. Choose the management type and read-only mode you want. Enter **<y>** 'for yes to allow the user to perform these operations, enter **<n>** for no if read-only mode is desired.

Using the HBAnyware Components

Note: To properly view the HBAnyware utility, ensure your system meets the following display requirements:
For Windows systems, the display resolution must be set to 800 by 600 or better.
For UNIX systems, the display resolution must be set to 1024 by 768 or better.
The display must run in 256-color mode or higher. HBAnyware icons use 256 colors. If the display is set for 16 color mode, HBAnyware icons are not displayed.

Starting the HBAnyware Utility

To start the HBAnyware utility:

In Windows:

On the Windows desktop, select **Start>All Programs>Emulex>HBAnyware**.

In Solaris LPFC, Solaris SFS and Linux:

1. Log on as 'root'.
2. Run the script to start the HBAnyware utility.
 - On Solaris LPFC and Solaris SFS:
`/opt/HBAnyware/hbanyware`
 - On Linux:
`/usr/sbin/hbanyware/hbanyware`

Starting the HBAnyware Utility with Web Launch

After the HBAnyware Web Launch software is installed and the Web Launch server is initialized, you can launch the HBAnyware utility directly with your Web browser.

Note: Only the HBAnyware Web Launch GUI is exported to the requesting client. All adapter discovery and remote management operations are performed by resources running on the remote host that served up the GUI component. Therefore, the SAN view displayed by the GUI is not from the client running the GUI, but rather from the host from which this GUI was retrieved.

To launch the HBAnyware utility with your Web browser:

1. Open your Web browser.
2. Enter the URL of an HBAnyware.jnlp file. Make sure that the URL specifies a remote server which has the HBAnyware Web Launch software installed and running.

`http://IP_ADDR:PORT_NUM/hbanyware.jnlp`

where *IP_ADDR* is the IP address of the host on which you installed the HBAnyware Web Launch Service, and *PORT_NUM* is the TCP port number of the listening hosts' Web server. The standard HBAnyware utility user interface is displayed.

The HBAware Utility Window Element Definitions

The HBAware utility window contains five basic components: the menu bar, the toolbar, the discovery-tree, the property tabs and the status bar.

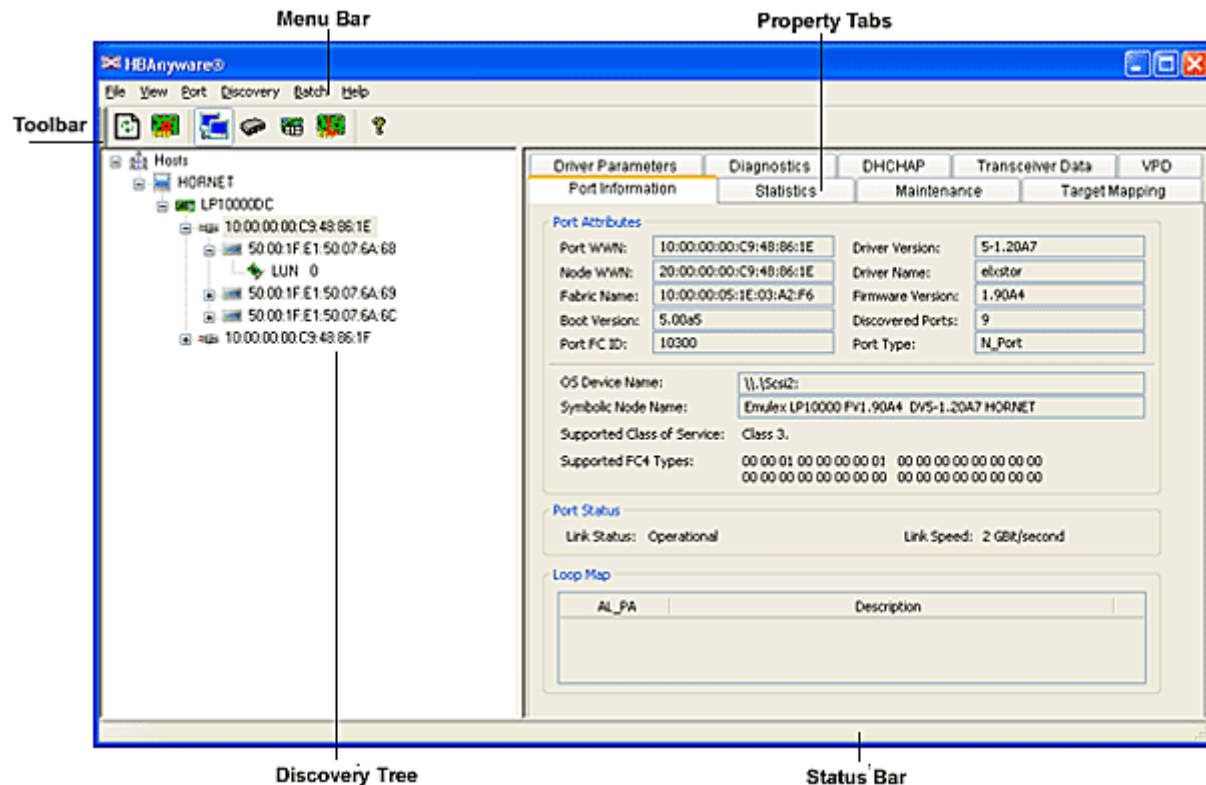


Figure 3: HBAware Utility window

Note: The element you select in the discovery-tree determines whether a menu item or toolbar icon is active. For example, if you select the local host or other system host, the Reset Adapter item on the Adapter menu is unavailable. The Reset Adapter toolbar button is unavailable as well.

Note: Screenshots in this manual are for illustrative purposes only. Your system information can vary.

Note: The features displayed by your local HBAware utility interface will match those of the remote server. When accessing a remote server on which an older version of the HBAware utility is installed, features that are not supported by the server's version of the HBAware utility will be unavailable.

The Menu Bar

The menu bar contains command menus that enable you to perform a variety of tasks such as exiting the HBAnyware utility, resetting adapters and sorting items in the discovery-tree view. Many of the menu bar commands are also available from the toolbar.

The Toolbar

The toolbar contains buttons that enable you to refresh the discovery-tree, reset the selected adapter and choose how you want to view discovered SAN elements in the discovery-tree. Many of the toolbar functions are also available from the menu bar.



Figure 4: Toolbar

The toolbar is visible by default. Use the Toolbar item in the View menu to hide the toolbar. If the item is checked, the toolbar is visible.

The Toolbar Buttons

The toolbar buttons perform the following tasks:



Click the **Discovery Refresh** button to force a full refresh cycle. A full refresh finds any new targets or virtual ports that were added to the SAN and removes any targets or virtual ports that were removed.



Click the **Reset** button to reset the selected adapter.

View Toolbar Buttons

The view toolbar buttons enable you to view SAN elements from the host, fabric, virtual ports, or by local or remote adapter perspective. By default, both local and remote adapters are displayed in Host view. The HBAnyware utility displays elements in ascending order.



Host View button (default)

- Displays the host system.

Note: You cannot change host names using the HBAnyware utility; names must be changed locally on that system.

- Within each host system, displays installed adapters.
- Displays adapter ports and port numbers if available.
- If multiple adapters have the same model number, displays adapters by WWNN.
- If targets are present, displays WWPN. Multiple adapters can refer to the same target.
- If LUNs are present, displays the LUN number.



Fabric View button

- Displays fabrics in the SAN with their fabric IDs.
- Displays ports under each switch.
- If targets are present, displays each WWPN. Multiple adapters can refer to the same target.
- If LUNs are present, displays each LUN number.
- If the fabric ID is all zeros, no fabric is attached.



Virtual Ports View button

- Displays virtual ports in the SAN.



Local HBAs Only button

- The Local HBAs Only menu item and button both work with the Host View and Fabric View. The first time you select this menu item or click this button, only local adapters are displayed. To change the view back to local and remote adapters, deselect the Local HBAs Only menu item or click the Local HBAs Only button again.



Help button

The Discovery-Tree

The discovery-tree (left pane) has icons that represent discovered hosts, adapters, ports, fabrics, targets and LUNs.

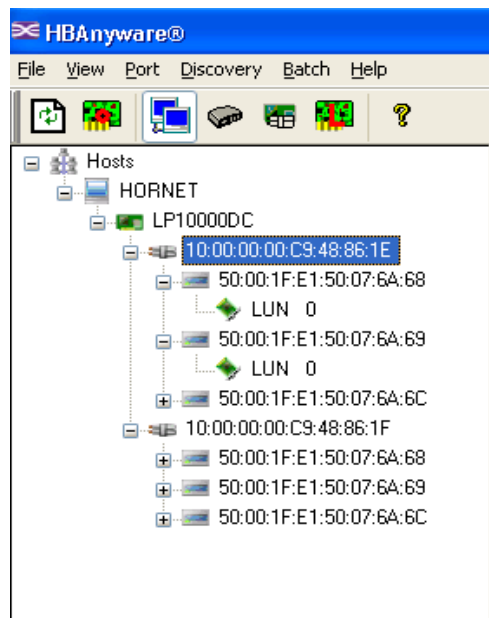


Figure 5: Discovery-tree

Discovery-Tree Icons

Discovery-tree icons represent the following:



The local host.



Other hosts connected to the system.



A green adapter icon with black descriptive text represents an online adapter.

A gray adapter icon indicates all ports for that adapter are no longer being discovered. A red icon indicates all ports for the adapter are offline (link down). Several situations could cause the adapter to be offline or inaccessible:

- The adapter on a local host is not connected to the network, but is available for local access.
- The adapter on a local host is malfunctioning and inaccessible to the local host and the network.
- The adapter on a local host is busy performing a local download and is temporarily inaccessible to the local host and the network.



The port icon represents adapter ports. Newer adapters also display the port number. A port icon with a red X indicates the port is down. If all discovered ports are down, the adapter icon changes to red. A gray port icon indicates that port is undiscovered. If all the ports are undiscovered, the adapter icon changes to gray.

Note: Multiport adapters are represented in the discovery-tree with separate port icons for each port. Older multiport adapter models (for example, LP8000DC, LP9402DC or LP9002DC) are represented by separate adapter icons.



The Virtual Port icon represents virtual ports.



The Target icon represents connections to individual storage devices.



The LUN icon represents connections to individual disk LUNs.



The Tape LUN icon represents LUNs that are tape devices.



The Target Controller LUN icon represents LUNs that are storage controllers.



The Switch icon represents connections to the switch.

Expanding or Collapsing the Discovery-Tree View

You can use the Expand/Collapse feature on the View menu to change the way discovered elements are displayed. By selecting one of the four levels, the discovery-tree is expanded or collapsed to that level. You can choose Hosts/Fabrics (depending on the view) HBAs, Ports and Targets.

Property Tabs

The property tabs display configuration, statistical and status information for network elements. The set of available tabs is context-sensitive, depending on the type of network element or adapter port currently selected in the discovery-tree.

Status Bar

The status bar is located near the bottom of the HBAnyware utility window. The status bar displays messages about certain HBAnyware utility functions, such as “Discovery in progress”.

The status bar is visible by default. Use the Status Bar item in the View menu to hide the status bar. If checked, the status bar is visible.

Customizing Tab Views

Using the Customize Tab Views dialog box you can choose whether or not to display certain property tabs. By default, all tabs are displayed.

To customize tab views:

1. From the **View** menu, select **Customize Tabs**. The Customize Tab Views dialog box appears.

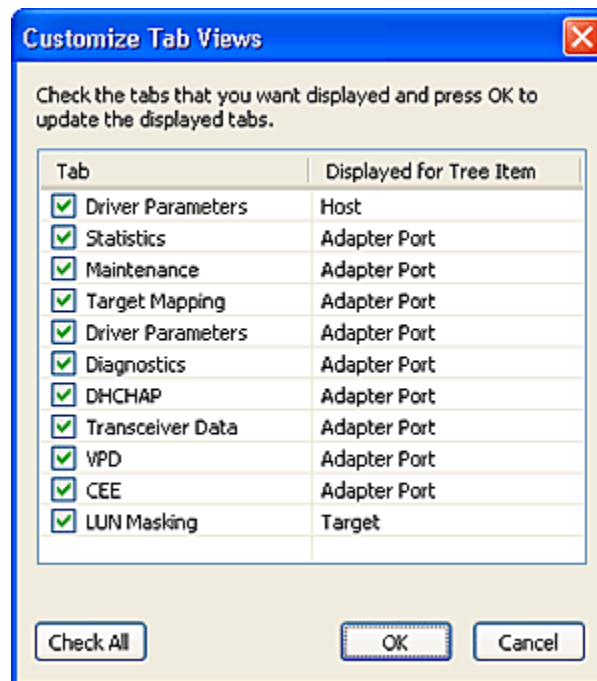


Figure 6: Customize Tab Views window

2. Check tabs to display them. Clear tabs to hide them.
3. Click **OK**.

Discovering HBAs

Automatic Fibre Channel Discovery

Adapters that have a physical FC connection to the same SAN are discovered automatically when the HBAnyware utility is launched. Adapters that don't have a physical FC connection to the SAN where the HBAnyware utility is launched can be discovered by sending management requests to a remote host using TCP/IP.

Note: The HBAnyware utility can only discover and manage remote adapters on hosts running the HBAnyware utility's remote management server. Remote FC capabilities of the HBAnyware utility are subject to fabric zoning and whether the HBAnyware utility's security is being used. Hosts you want to discover and manage using the HBAnyware utility must be in the same zone or discovered and managed through TCP/IP access.

Note: After adding an adapter to a running system (commonly called a hot plug), click **Discovery Refresh** (🔄) or restart the HBAnyware utility to display the new adapter port in the discovery-tree.

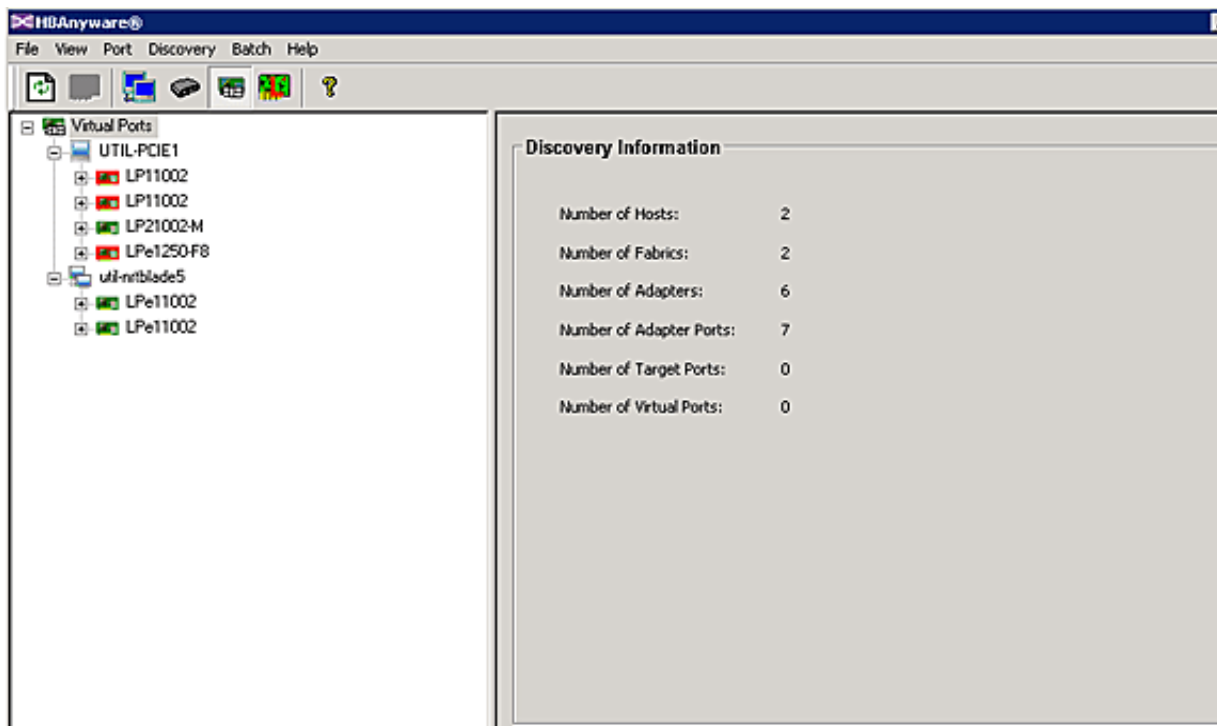


Figure 7: Discovery Information

Changing Adapter Port Names

The HBAnyware utility enables you to change adapter port names. (Not available in read-only mode.)

For example, you may want to identify a particular adapter port with the function it supports, such as a tape drive, scanner, or some other device. Use any characters you want for names, and names can be up to 255 characters in length. You can also revert to the adapter's default name.

Note: Although you can change the adapter port's displayed name from the default WWN, the change occurs in the discovery-tree only. The WWN is still active, it is simply replaced for display purposes with the name you enter. For example, the Port WWN field of the Port Information tab is not changed. Also, any change you make to the adapter port names in your discovery-tree are seen only by you; users running the HBAnyware utility on another host do not see your name changes.

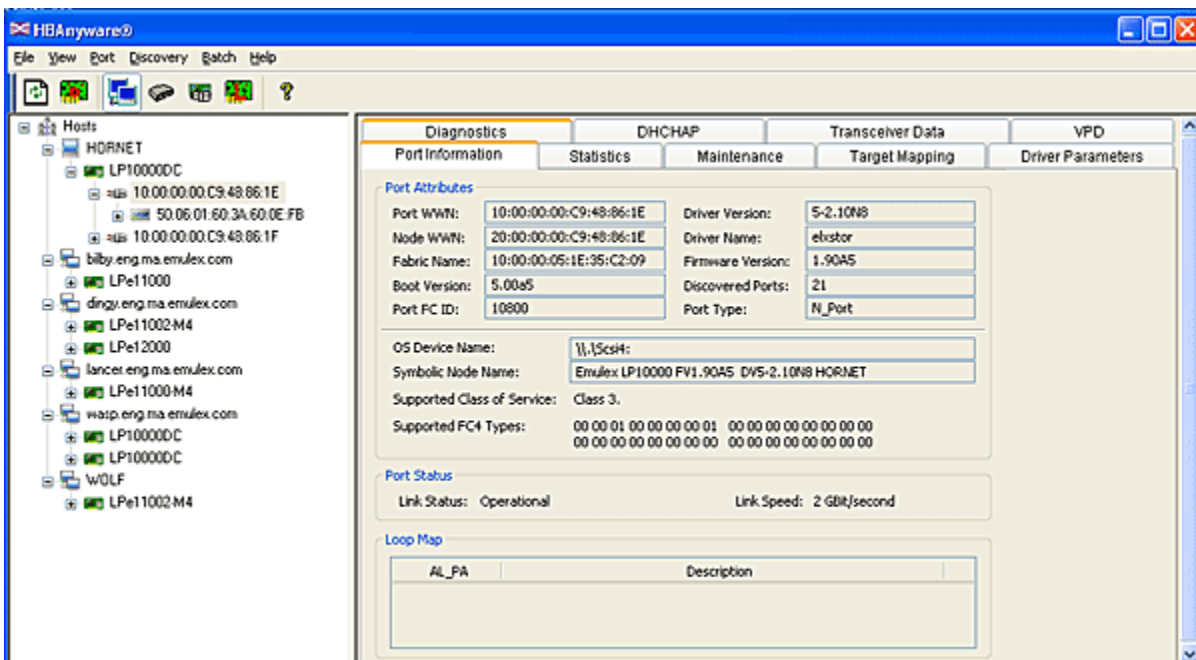


Figure 8: Port Information tab

To change the name of an adapter:

1. From the discovery-tree, select the port whose name you want to change.
2. Do one of the following:
 - Select **Edit Name** from the **Port** menu.
 - From the discovery-tree, right-click the port whose name you want to change and select **Edit Name** (or **Change Name**).
3. Edit the port name in the discovery-tree.

To use the adapter port's default name:

1. From the discovery-tree, select the port whose name you want to change.
2. Do one of the following:
 - Select **Use Default Name** from the **Port** menu.
 - From the discovery-tree, right-click the port whose name you want to change and select **Use Default Name** (or **Restore Default Name**).

Remote SAN Management Using TCP/IP Access Protocol

You can discover adapters on TCP/IP hosts. Remote SAN management over TCP/IP sends remote management requests using TCP/IP access protocol to remote hosts. TCP/IP access enables you to access adapters via their host IP-address or by the name of the host on which they reside. Since adapters can exist on a host but not be a part of a FC network, they do not appear during normal FC discovery. Thus, TCP/IP access enlarges the number of adapters that can be queried or modified.

Note: In Windows, if you are running a firewall you may need to add the HBAnyware utility remote server to the firewall's exception list. This remote server's path is:
 \Program Files\Emulex\Util\Common\rmserver.exe
On 64-bit hosts the path is
 \Program Files (x86)\Emulex\Util\Common\rmserver.exe

The principle differences between FC and TCP/IP access are:

- A TCP/IP host with or without an adapter installed does not need to connect to a fabric to manage other hosts.
- A TCP/IP management host can manage all of the adapters in a remote host, not just the ones connected to the same fabric. FC can only manage adapters connected to the same fabric.
- You can manage many more hosts since TCP/IP access is not constrained by the boundaries of a fabric or zoning.
- True board status (e.g. link down) is available since the FC path is not necessary to send a status request to the remote host.
- Adapter security in a TCP/IP environment is much more important since many more hosts are available for management and TCP/IP access is not affected by fabrics or zoning.
- Discovery of hosts in a TCP/IP environment is not automatic as FC discovery is. You must add the hosts to be managed.

The Hosts File

The TCP/IP discovery portion of the HBAnyware utility discovery server relies on a file called the hosts file. This plain text file contains a list of hosts the utility will attempt to discover. The discovery server does not attempt to discover hosts over TCP/IP through any other mechanisms (e.g. ping sweeps, broadcasts, etc.).

The hosts file is automatically created or modified when you perform any of the following operations:

- Add a single host from the Add Remote Host window. If the host is discovered, the HBAnyware utility adds its IP address and name to the host file.
- Scan a range or ranges of IP addresses for hosts that can be managed. This is performed in the Add Remote Hosts window. For each discovered host, the HBAnyware utility adds the IP address and name to the host file.
- Remove a host from the host file from the Remove Remote Hosts window. For each removed host, the HBAnyware utility removes that IP address and name from the host file. Manually edit the file to add or remove hosts.

Adding a Single Host

The HBAnyware utility enables you to specify a single TCP/IP host to manage. If the host is successfully discovered it is added to the hosts file. If it has not been discovered over FC already, the host and its adapter ports are added to the discovery-tree. (Not available in read-only mode.)

Prerequisites

The HBAnyware utility must be installed on the remote host.

Procedure

To add a single host:

1. From the **Discovery** menu, select **TCP/IP>Add Host**. The Add Remote Host dialog box appears.

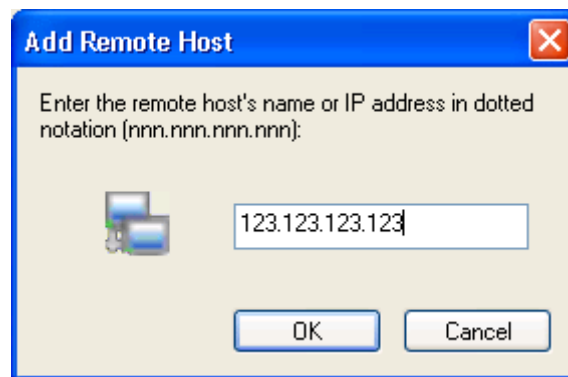


Figure 9: Add Remote Host dialog box

2. Enter the name or the IP address of the host to be added.

Note: Entering the IP address to identify the host avoids name resolution issues.

3. Click **OK**. You will receive a message indicating whether the new host was successfully added.

Adding a Range of Hosts

Find the TCP/IP-accessed manageable hosts by searching a range of IP addresses. The Add Remote Hosts dialog box enables you to build the initial list of TCP/IP accessed manageable hosts. (Not available in read-only mode or on Windows XP or Vista.)

Note: The ranges of IP addresses are only scanned each time you open the Add Remote Hosts dialog box and click Start. The ranges are NOT automatically scanned by the discovery server during its discovery cycles.

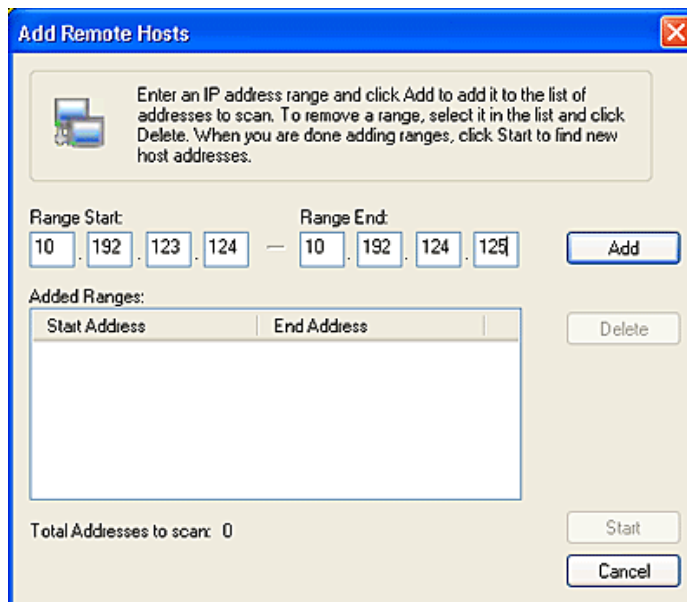


Figure 10: Add Remote Hosts dialog box

Prerequisites

The HBAware utility must be installed on all remote hosts.

Procedure

To add a range of remote hosts:

1. From the **Discovery** menu, select **TCP/IP>Add Remote Hosts**. The Add Remote Hosts dialog box appears.
2. Enter the complete start and end address range and click **Add**. The added address range appears in the dialog box. Add any additional ranges you want to search.

3. Click **Start**. If an address is determined to be remotely manageable it is added to the list of addresses the discovery server will attempt to discover. The utility creates a hosts file if necessary, and checks each address in the range to determine if the host is available and remotely manageable. The number of addresses (of manageable hosts) discovered is periodically updated on the dialog box.

Note: The number of addresses does not correspond directly to the number of hosts added to the discovery-tree.

For example, some of the addresses discovered may be for hosts that have already been discovered over FC. However, new adapters can be discovered on those hosts that were not discovered over FC.

Also, a host can have more than one adapter installed and both IP addresses for that host are discovered during the search, but only one host can possibly be added to the discovery-tree.

4. Save the IP address ranges:
 - **In Windows:** A dialog box appears asking you to save the IP address ranges you searched. Click **Yes** to save the address ranges. If you save the address ranges, these address ranges appear the next time you use the Add Range of IP Hosts dialog box. Click **No** if you do not want to save the address ranges.
 - **In Solaris LPFC, Solaris SFS and Linux:** Click **Save Ranges to File** to save the specified range(s) to a file so that these address ranges appear the next time you use the Add Range of IP Hosts dialog box.

Removing Hosts

Removing hosts that can no longer be discovered improves the operation of the discovery server. For example, you may want to remove a host when it is removed from the network. (Not available in read-only mode.)

To remove hosts:

1. From the **Discovery** menu, select **TCP/IP>Remove Host(s)**. The Remove Remote TCP/IP Hosts dialog box shows a list of discovered hosts. Any host not currently discovered appears in red. Click **Show Undiscovered Hosts Only** to display only currently undiscovered hosts.
2. From the Remove Remote TCP/IP Hosts dialog box, select the hosts you want to remove. You can select all the displayed hosts by clicking **Select All**.
3. Click **OK** (or **Remove**) to remove the selected hosts.

Manually Editing the Hosts File

You can open the hosts file with any text editor, modify the contents and save the file. The name of the host file is "hbahosts.lst". Once the file is modified and saved, the updated file is used after the next TCP/IP discovery cycle is complete. If the discovery server is running, it does not need to be restarted.

To manually edit the hosts file:

1. Locate and open the hosts file.

Windows: The file is located on the system drive in the directory "\\Program Files\\Emulex\\Util" for 32-bit machines or "\\Program Files (x86)\\Emulex\\Util" for 64-bit machines.

Solaris: The file is located in the directory "/opt/HBAnyware".

Linux: The file is located in the directory "/usr/sbin/hbanyware".

2. Edit the file. Guidelines for editing the file are as follows:
 - Each line of the file starts with an IP address. Following the IP address can be any number of tabs or spaces. This is followed by a “#” character, zero or more tabs or spaces and the name of the host for that IP address. The host name is not required for discovery. Its purpose is to make the file more readable and is used by the HBAnyware utility to display the host name in the Remove Remote Hosts window when the host is not discovered. However, the discovery server only needs the IP address to discover the host.
 - Each line in the file can be up to 1023 characters, although this is longer than is needed for a host IP address and host name. A line longer than this is truncated, possibly causing discovery to not discover some of the hosts.
 - Blank lines are ignored.
3. Save the file.

Copying the File

A hosts file on one host can be copied and used on another host. This is useful if there are multiple hosts on the same network running the HBAnyware utility. Once the remote hosts are added to the hosts file on one host, that hosts file can be copied to other hosts so the process to create the hosts file does not need to be repeated.

Note: Due to the line terminator differences between Windows and Solaris or Linux hosts, the files cannot be shared between Windows hosts and Solaris or UNIX hosts.

Configuring Discovery and TCP/IP Settings

Use the Modify Discovery Settings dialog box to configure several discovery server parameters. You can define when to start the discovery server, when to refresh FC and TCP/IP accessed discoveries and when to remove previously discovered HBAs that are no longer being discovered. For TCP/IP management, you can specify an IP port number, change an IP port number and enable a port for TCP/IP management.

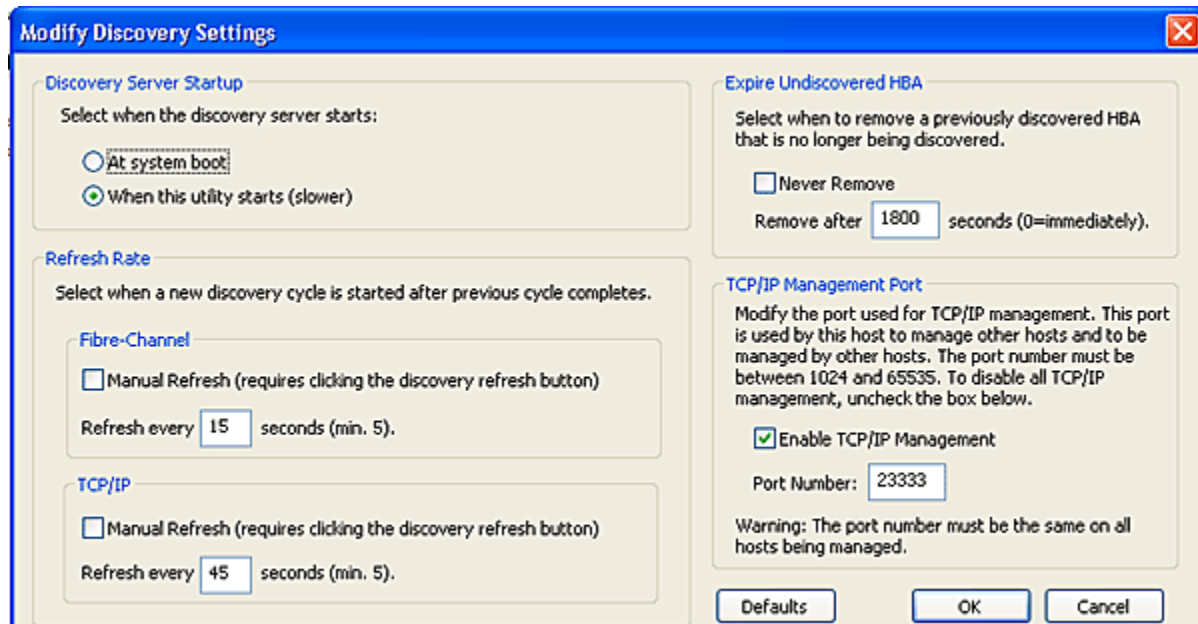


Figure 11: Adapter Discovery Properties dialog box

To configure discovery settings:

1. From the **Discovery** menu, select **Modify Settings**. The Modify Discovery Settings dialog box appears.
2. Define the discovery properties you want and click **OK**. Click **Defaults** to return the discovery properties to their default settings.
3. If TCP/IP Management is enabled, the Enable TCP/IP Management checkbox is selected and the current port number is displayed in the Port Number field. If desired, enter a different port number (between 1024 and 65535). Click **Defaults** to select the Enable TCP/IP Management checkbox (if unchecked) and set the port number to 23333.

If the port number or the Enable TCP/IP Management checkbox is changed, a set of warning messages may appear before changes are made. Click **Yes** on the warning message to continue with the change.

Caution: The IP port number must be the same for all hosts that are to be managed. Setting an IP port number for one host to a different value than the other hosts will make the host unable to manage other hosts over TCP/IP, as well as make the host unmanageable over TCP/IP from other hosts.


4. If the IP port number is changed, the utility restarts the HBAnyware utility discovery server and management agent to use the new settings. If the servers cannot be stopped and restarted, you are prompted to reboot the host for the new TCP/IP management settings to take effect.

Resetting Adapter Ports

You can reset remote and local adapter ports. (Not available in read-only mode).

Caution: Do not reset your adapter port while copying or writing files. This could result in data loss or corruption.

To reset the adapter port:

1. In the discovery-tree, select the adapter port you want to reset.
2. Do one of the following:
 - From the **Port** menu, click **Reset Adapter**.
 - Click the **Reset** toolbar button: .

The following warning appears:

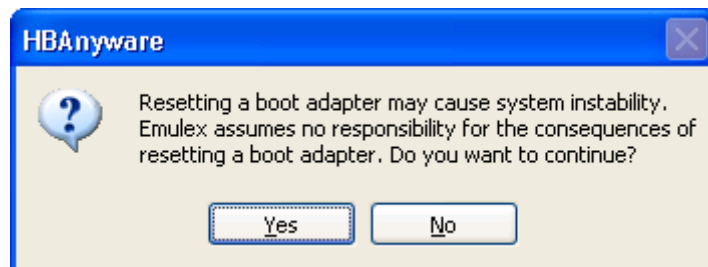


Figure 12: Reset Warning


3. Click **Yes**. The adapter port resets.

The reset can require several seconds to complete. While the adapter port is resetting, the status bar shows "Reset in progress." When the reset is finished, the status bar shows "Reset Completed".

Viewing Host Information

There are two tabs that show host information: the Host Information tab and the Driver Parameters tab. The Host Information tab is read-only. The Driver Parameters tab enables you to view and define adapter driver settings for a specific host.

To view the Host Information and Driver Parameters tabs:

1. Do one of the following:
 - From **View** menu, click **Hosts**.
 - From the toolbar, click  **Host View**.
2. Select a host in the discovery-tree.
3. Select the **Host Information** tab or the **Driver Parameters** tab.

The Host Information Tab

The Host Information tab displays information for the selected host including the number of adapters installed in the selected host, the number of fabrics to which it is connected and so on.

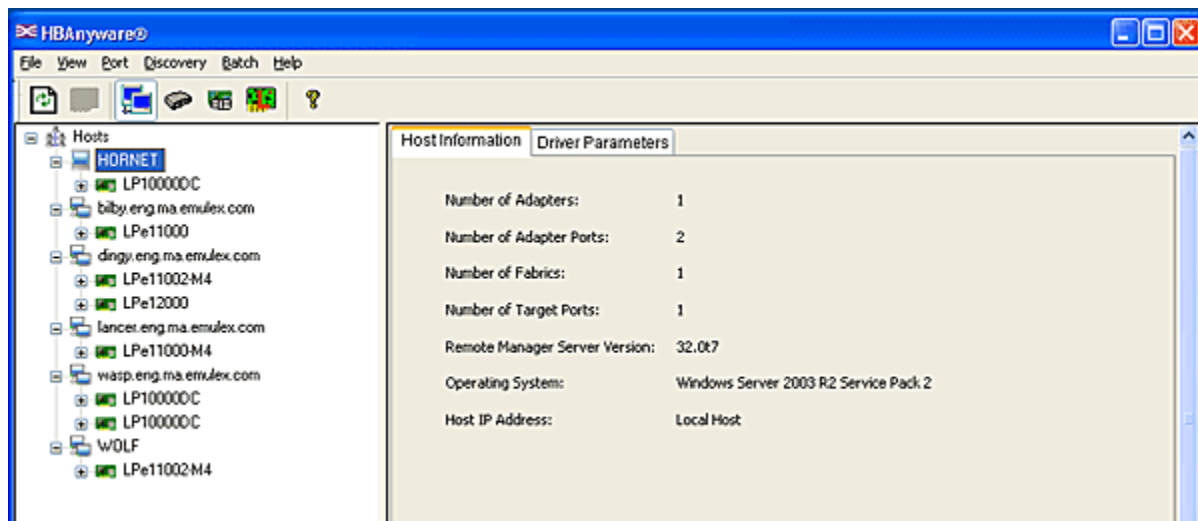


Figure 13: Host Information tab

Host Information Field Definitions

- Number of Adapters - The number of adapters installed in the host.
- Number of Adapter Ports - The number of discovered adapter ports on this host that can be managed by this host.
- Number of Fabrics - The number of fabrics to which this host is attached.
- Number of Target Ports - The number of storage devices seen by the host.
- Remote Manager Server Version - The version of the HBAnyware utility server that is running on the host. If different versions of the HBAnyware utility are installed on different hosts in the SAN, those differences appear in this field.
- Operating System - The operating system and version installed on the selected host.
- Host IP Address - If the host is discovered with FC, the dialog box displays "Host discovered over Fibre Channel". If the host has been added with TCP/IP access, the Host IP Address field displays the host's IP address, e.g., 138.239.82.131.

The Driver Parameters Tab

The Driver Parameters tab enables you to view and edit the adapter driver parameter settings contained in a specific host. The host driver parameters are global values and apply to all adapters in that host unless they are overridden by parameters assigned to a specific adapter using the adapter Driver Parameters tab. For each parameter, the tab shows the current value, the range of acceptable values, the default value, and whether the parameter is dynamic. A dynamic parameter allows the change to take effect without resetting the adapter or rebooting the system.

Note: For the Linux 2.6 kernel, most driver parameters are set globally. You can set the `lpfc_log_verbose`, `lpfc_nodev_tmo` and `lpfc_use_adisc` locally.

Note: For all compatible Linux versions: If you change driver parameters using the HBAnyware utility and you want these changes to be permanent and persist across system reboots, you must create a new ramdisk image. The ramdisk image is used when the kernel is initialized during system startup, and loads the LPFC driver with the updated driver parameters.

To create a new ramdisk you can use the LPFC driver's `lpfc-install` script. Refer to the “Creating a New Ramdisk” section of the Emulex Driver for Linux User Manual for instructions.

For information on changing parameters for a single adapter, see “Setting Driver Parameters” on page 54. For information on changing parameters for the host, see “Setting Driver Parameters for All HBAs in a Host” on page 56.

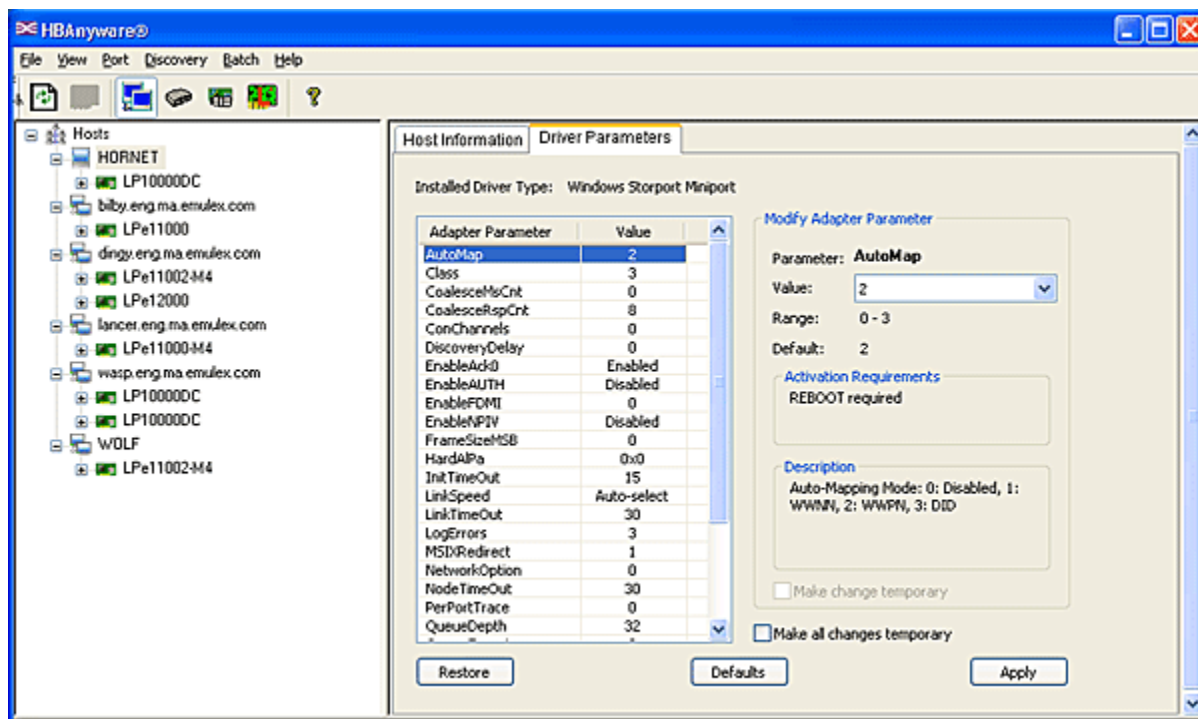


Figure 14: Driver Parameters tab

Note: If there is more than one driver type installed, the Installed Driver Types menu shows a list of all driver types and driver versions that are installed on the adapters in the host.

Driver Parameters Tab Field Definitions

- Installed Driver Type - The current driver and version installed on this host.
- Adapter Parameter table - A list of adapter driver parameters and their current values.

Modify Adapter Parameter Area


- Adapter-specific information displays in this area. This can include value, range, default, activation requirements and description.

Driver Parameters Tab Buttons (Not available in read-only mode.)

- Restore - If you changed driver parameters, but did not click **Apply** and you want to restore the parameters to their last saved values, click **Restore**.
- Defaults - Click to reset all driver parameter values to their default (out-of-box) values.
- Apply - Click to apply any driver parameter changes. If you changed a driver parameter that is not dynamic, you may need to reset the adapter port or reboot the system.

Viewing Fabric Information

The Fabric Information tab displays information about the selected fabric.

1. Do one of the following:
 - From the **View** menu, select **Fabric**.
 - From the toolbar, click  **Fabric View**.
2. Click on a fabric address in the discovery-tree. The Fabric Information tab shows information about the selected fabric.

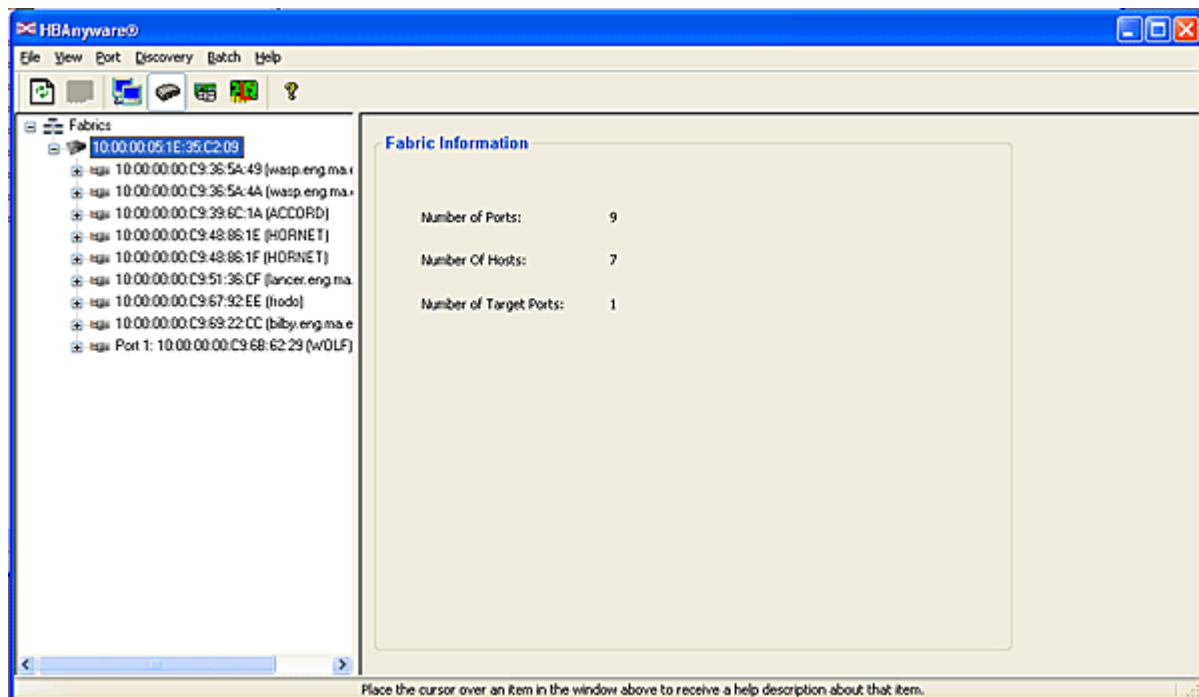


Figure 15: Fabric Information


Fabric Information Field Definitions

- Number of Ports - The number of discovered adapter ports on this host that can be managed by this host.
- Number of Hosts - The number of hosts discovered or seen by this host on the selected fabric.
- Number of Target Ports - The number of storage devices seen by this host on the selected fabric.

Viewing Virtual Port Information

View virtual port information and their associated targets and LUNs.

To view virtual port information:

1. Do one of the following:
 - From the **View** menu, select **Virtual Ports**.
 - From the toolbar, click  **Virtual Ports View**.

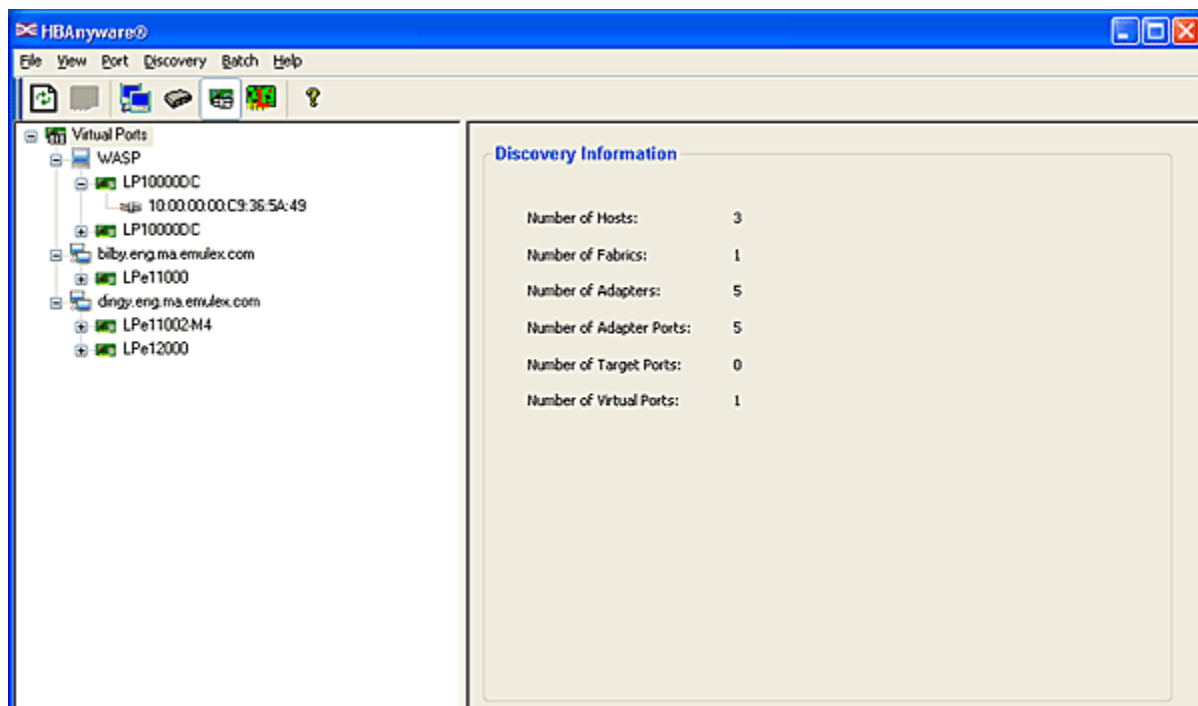


Figure 16: Virtual Ports Discovery Information

Virtual Port Information Field Definitions

- Number of Hosts - The total number of hosts discovered in the SAN.
- Number of Fabrics - The total number of fabrics discovered in the SAN.
- Number of Adapters - The total number of adapters discovered in the SAN.
- Number of Adapter Ports - The total number of adapter ports discovered in the SAN.
- Number of Target Ports - The total number of target ports discovered in the SAN.
- Number of Virtual Ports - The total number of virtual ports discovered in the SAN.

Viewing Discovery Information

Discovery Information contains a general summary of the discovered elements. The Host, Fabric or Virtual Port icon, depending upon which view you select, is the root of the discovery-tree, but it does not represent a specific network element. Expanding it reveals all hosts, LUNs, targets, adapters and ports that are visible on the SAN.

To view the discovery information:

1. Click the **Hosts**, **Fabrics** or **Virtual port** icon at the root of the discovery-tree. Discovered SAN elements appear in the discovery-tree.
2. Select an element from the discovery-tree to learn more about it.

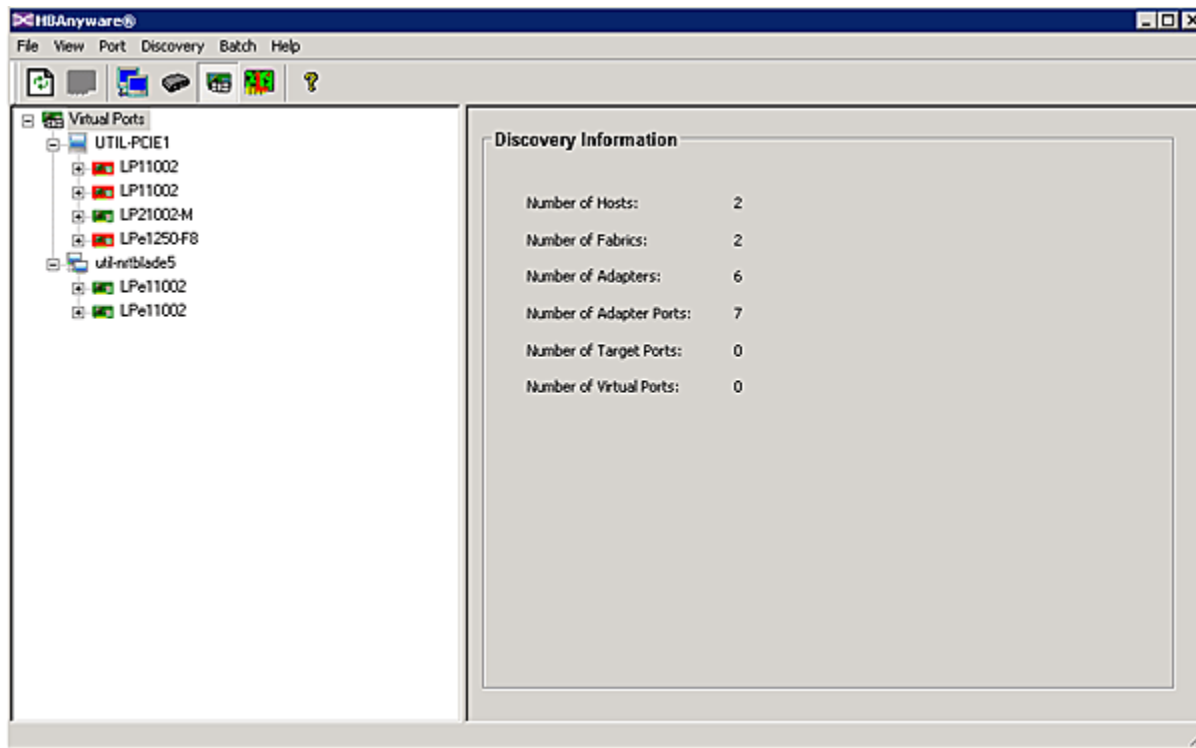


Figure 17: Discovery Information

Discovery Information Field Definitions

- **Number of Hosts** - The total number of discovered host computers. This includes servers, workstations, personal computers, multiprocessor systems and clustered computer complexes.
- **Number of Fabrics** - The total number of discovered fabrics.
- **Number of Adapters** - The total number of discovered adapters.
- **Number of Adapter Ports** - The number of discovered adapter ports on this host that can be managed by this host.
- **Number of Target Ports** - The total number of unique discovered targets on the SAN. In the discovery-tree, the same target can appear under more than one adapter.
- **Number of Virtual Ports** - The number of discovered virtual ports on this host that can be managed by this host. (Only displayed if the Virtual Ports element is selected in the discovery-tree.)

Viewing Adapter Information

The Adapter Information tab contains general attributes associated with the selected adapter.

To view general adapter information:

1. Select **Host View** or **Virtual Ports View**.
2. Select an adapter in the discovery-tree.

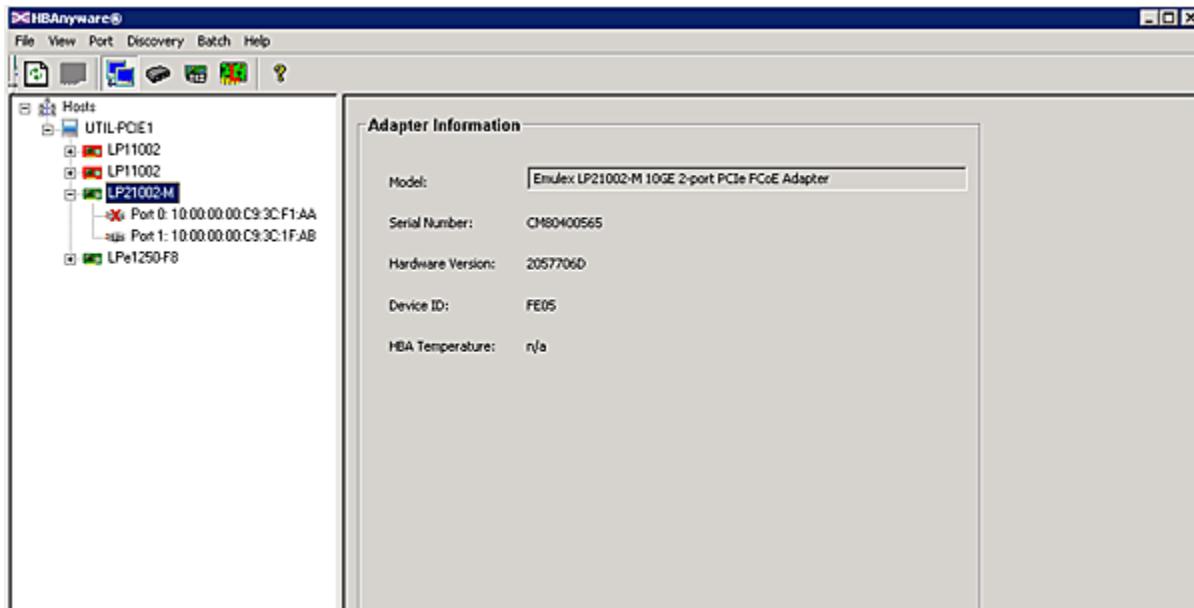


Figure 18: Adapter Information tab

Adapter Information Field Definitions

- **Model** - The complete model name of the adapter.
- **Serial Number** - The manufacturer's serial number for the selected adapter.
- **Hardware Version** - The board JEDEC ID version for the selected adapter.
- **Device ID** - The default device ID for the selected adapter.
- **HBA Temperature** - If the adapter's temperature is not available, "n/a" is displayed. If supported by the adapter, this field displays the adapter's temperature and one of the following temperature-related status messages:
 - **Normal**: The adapter's temperature is within normal operational range.
 - **Exceeded operational range - Critical**: The adapter's temperature is beyond normal operational range. If the temperature continues to increase, the adapter shuts down. You must determine the cause of the temperature problem and fix it immediately. Check for system cooling issues. Common causes of system cooling issues include clogged air filters, inoperable fans and air conditioning problems that cause high ambient air temperatures.
 - **Exceeded operational range - Adapter stopped**: The temperature has reached critical limit, forcing the adapter to shut down. You must determine the cause of the temperature problem and fix it before resuming operation. Check for system cooling issues. Common causes of system cooling issues include clogged air filters, inoperable fans and air conditioning problems that cause high ambient air temperatures.

After the system overheating issue is resolved and the adapter has cooled down, reboot the system or, if the system supports hot swapping, cycle the power of the adapter slot.

Viewing Port Information

The Port Information tab contains detailed information associated with the selected adapter port.

To view port information:

1. Select **Host View** or **Fabric View**.
2. Select an adapter port in the discovery-tree.
3. Select the **Port Information** tab.

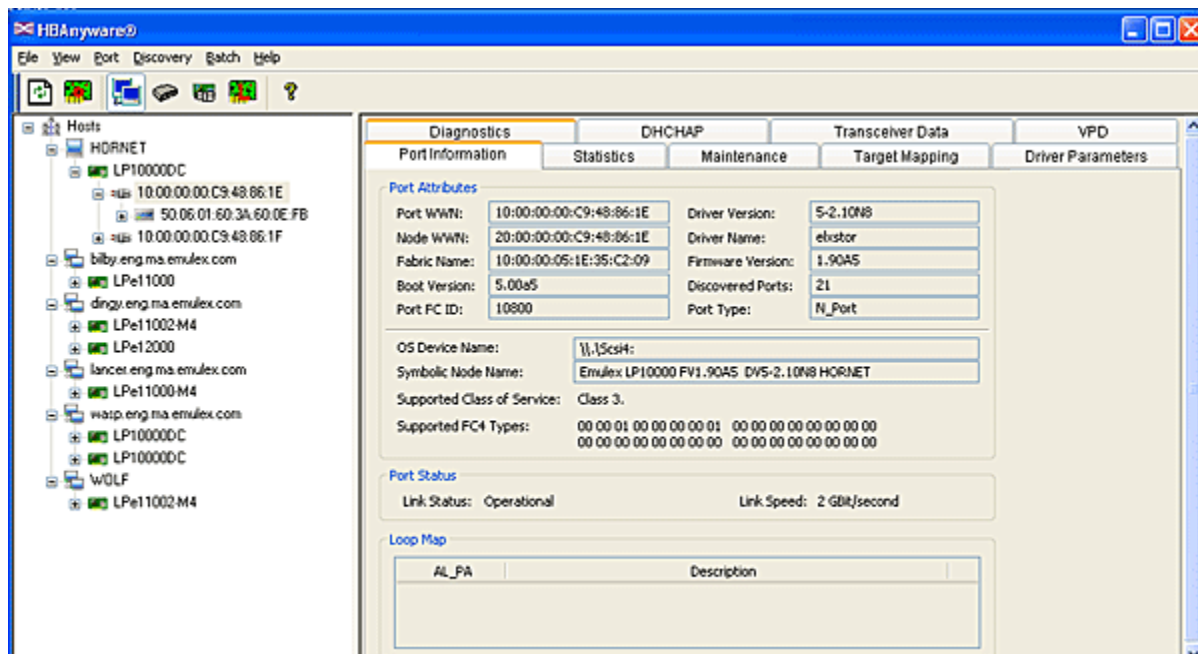


Figure 19: Port Information tab

Port Attributes Area Field Definitions

- Port WWN - The Port World Wide Name of the adapter.
- Node WWN - The Node World Wide Name of the selected adapter.
- Fabric Name or Host Name - The Fabric Name field is displayed in Host view. This is a 64-bit worldwide unique identifier assigned to the fabric. The Host Name is displayed in Fabric view. The host name is the name of the host containing the adapter.
- Boot Version - The version of boot code installed on the selected adapter port. If the boot code is disabled, the field displays "Disabled".
- Port FC ID - The Fibre Channel ID for the selected adapter port.
- Driver Version - The version of the driver installed for the adapter.
- Driver Name - The executable file image name for the driver as it appears in the Emulex driver download package.
- Firmware Version - The version of Emulex firmware currently active on the adapter port.
- Discovered Ports - Counts the number of mapped and unmapped ports found during discovery by the Emulex adapter driver. The mapped ports are targets and the unmapped ports are non targets such as switches or adapters.
- Port Type - The current operational mode of the selected adapter's port.

- OS Device Name - The platform-specific name by which the selected adapters is known to the operating system (OS).
- Symbolic Node Name - The FC name used to register the driver with the name server.
- Supported Class of Service - A frame delivery scheme exhibiting a set of delivery characteristics and attributes. There are three classes of service.
 - Class-1 provides a dedicated connection between a pair of ports with confirmed delivery or notification of non-delivery.
 - Class-2 provides a frame switched service with confirmed delivery or notification of non-delivery.
 - Class-3 provides a frame switched service similar to Class-2 but without notification of frame delivery or non-delivery.
- Supported FC4 Types - A 256-bit (8-word) map of the FC-4 protocol types supported by the port containing the selected adapter.

Port Status Area Field Definitions

- Link Status - The status of the link on the selected adapter port.
- Link Speed - The current link speed of the selected adapter port.

Loop Map Table Definitions

- The loop map shows the different ports present in the loop, and is present only if the port (adapter) is operating in loop mode. The simplest example would be to connect a JBOD directly to an adapter. When this is done, the port type is a private loop, and the loop map has an entry for the adapter, and one entry for each of the disks in the JBOD.

Viewing Port Statistics

The Statistics tab provides cumulative totals for various error events and statistics on the port. Some statistics are cleared when the adapter is reset.

To view port statistics:

1. Select **Host View** or **Fabric View**.
2. Select an adapter port in the discovery-tree.
3. Click the **Statistics** tab.

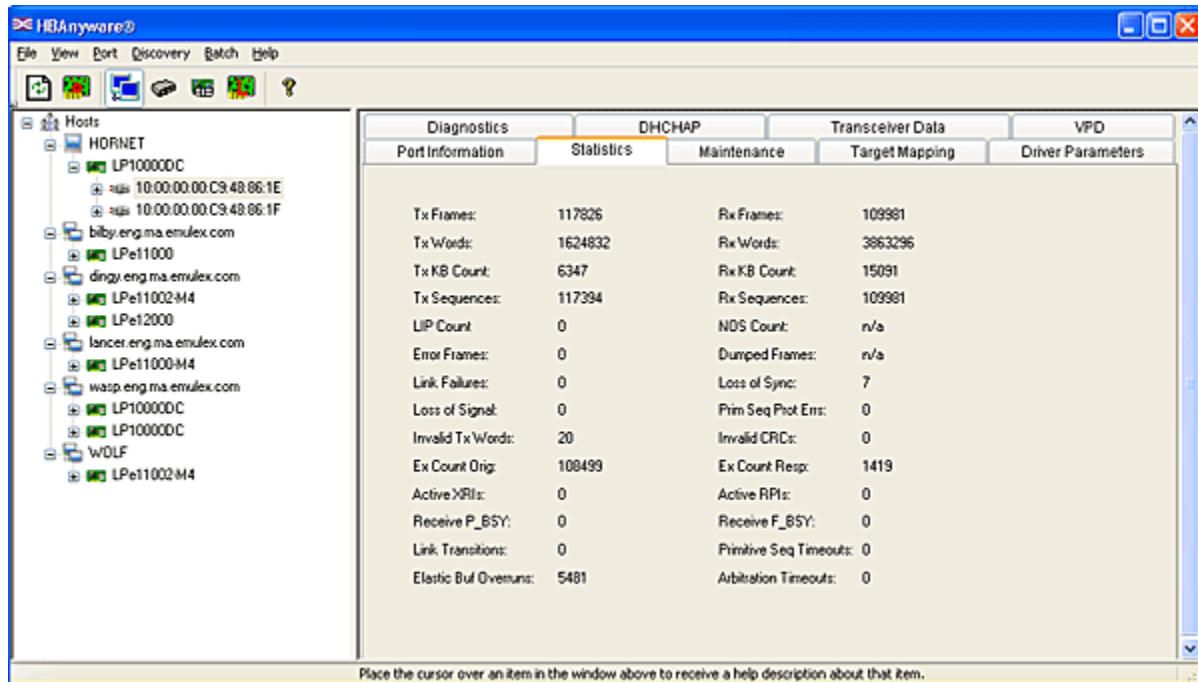


Figure 20: Statistics tab

Port Statistics Field Definitions


- Tx Frames - FC frames transmitted by this adapter port.
- Tx Words - FC words transmitted by this adapter port.
- Tx KB Count - FC kilobytes transmitted by this adapter port.
- Tx Sequences - FC sequences transmitted by this adapter port.
- LIP count - The number of loop initialization primitive (LIP) events that have occurred for the port. This field is not supported if the topology is not arbitrated loop. Loop initialization consists of the following:
 - Temporarily suspending loop operations.
 - Determining whether loop capable ports are connected to the loop.
 - Assigning AL_PA IDs.
 - Providing notification of configuration changes and loop failures.
 - Placing loop ports in the monitoring state.
- Error Frames - The number of frames received with cyclic redundancy check (CRC) errors.
- Link Failures - The number of times the link failed. A link failure is a possible cause of a timeout.

- Loss of Signal - The number of times the signal was lost.
- Invalid Tx Words - The total number of invalid words transmitted by this adapter port.
- Ex Count Orig - The number of FC exchanges originating on this port.
- Active XRIs - The number of active exchange resource indicators.
- Received P_BSY - The number of FC port-busy link response frames received.
- Link Transitions - The number of times the SLI port sent a link attention condition.
- Elastic Buf Overruns - The number of times the link interface has had its elastic buffer overrun.
- Rx Frames - The number of FC frames received by this adapter port.
- Rx Words - The number of FC words received by this adapter port.
- Rx KB Count - The received kilobyte count by this adapter port.
- Rx Sequences - The number of FC sequences received by this adapter port.
- NOS count - This statistic is currently not supported for the SCSIport Miniport and Storport Miniport drivers, nor is it supported for arbitrated loop.
- Dumped Frames - This statistic is not currently supported for the SCSIport Miniport driver, the Storport Miniport driver or the driver for Solaris.
- Loss of Sync - The number of times loss of synchronization has occurred.
- Prim Seq Prot Errs - The primitive sequence protocol error count. This counter is incremented whenever there is any type of protocol error.
- Invalid CRCs - The number of frames received that contain CRC failures.
- Ex Count Resp - The number of FC exchange responses made by this port.
- Active RPIs - The number of remote port indicators.
- Receive F_BSY - The number of FC port-busy link response frames received.
- Primitive Seq Timeouts - The number of times a primitive sequence event timed out.
- Arbitration Timeouts - The number of times the arbitration loop has timed out. Large counts could indicate a malfunction somewhere in the loop or heavy usage of the loop.

Viewing Fabric Discovery Information

The Discovery Information tab contains information about the selected fabric.

To view fabric discovery information:

1. Do one of the following:
 - From the **View** menu, select **Fabric**.
 - From the toolbar, click  **Fabric View**.

The Discovery Information tab shows information about the fabric.

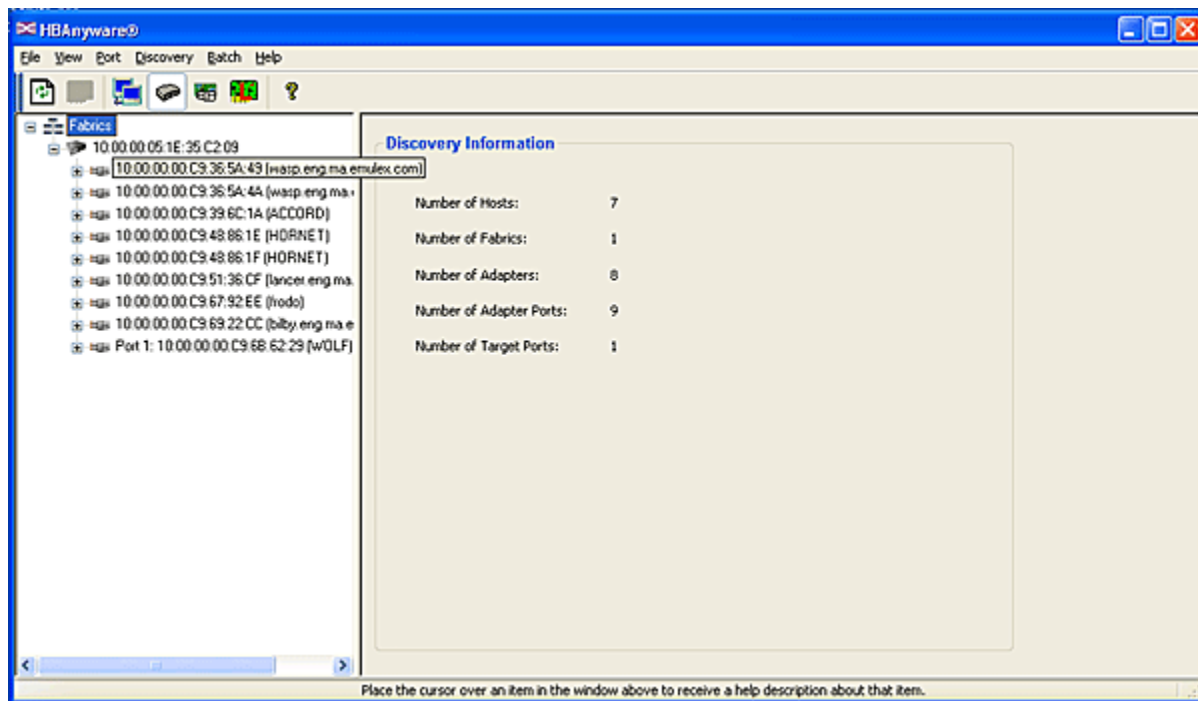


Figure 21: Fabric Discovery Information

Discovery Information Field Definitions

- Number of Hosts - The number of hosts discovered or seen by this host on the selected fabric.
- Number of Fabrics - The number fabrics identified during discovery.
- Number of Adapters - The number of adapters discovered by this host on the selected fabric.
- Number of Adapter Ports - The number of discovered adapter ports on this host that can be managed by this host.
- Number of Target Ports - The number of storage devices seen by this host on the selected fabric.

Viewing Transceiver Information

The Transceiver Data tab enables you to view transceiver information such as vendor name, serial number, part number and so on.

To view transceiver information:

1. Select **Host View** or **Fabric View**.
2. In the discovery tree, select the port whose transceiver information you want to view.
3. Select the **Transceiver Data** tab.

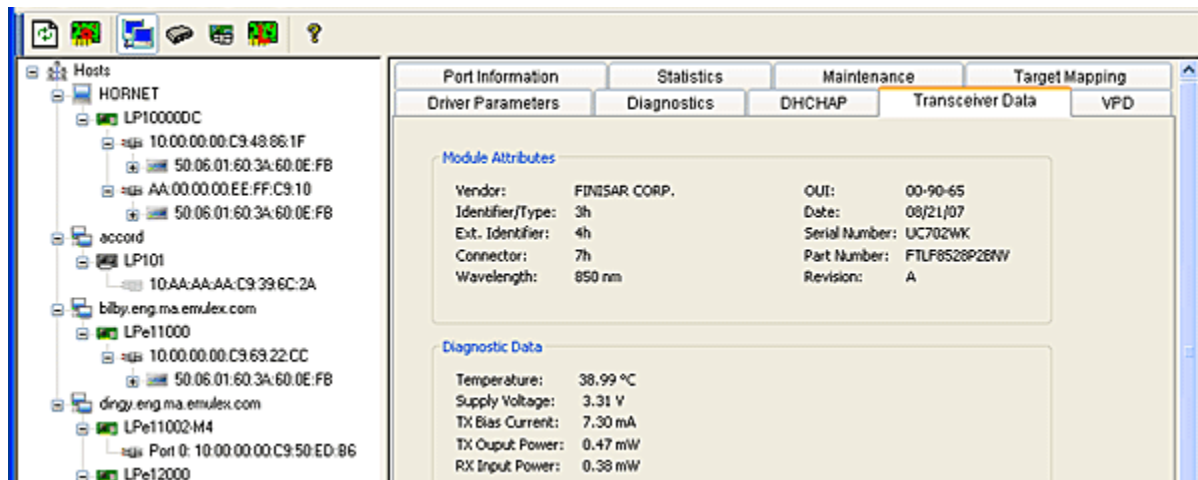


Figure 22: Transceiver Data tab

Transceiver Information Field Definitions

Module Attributes Area

- Vendor - The name of the vendor.
- Identifier/Type - The identifier value that specifies the physical device described by the serial information.
- Ext. Identifier - Displays additional information about the transceiver.
- Connector - The external optical or electrical cable connector provided as the media interface.
- Wavelength - The nominal transmitter output wavelength at room temperature.
- OUI - Displays the vendor Organizationally Unique Identifier. It is also known as the IEEE Company Identifier for the vendor.
- Date - The vendor's date code in the MM/DD/YY format.
- Serial Number - The serial number provided by the vendor.
- Part Number - The part number provided by the SFP vendor.
- Revision - The vendor revision level.

Diagnostic Data Area

- Temperature - The internally measured module temperature.
- Supply Voltage - The internally measured supply voltage in the transceiver.
- TX Bias Current - The internally measured TX Bias Current.
- TX Output Power - The measured TX output power
- RX Output Power - The measured RX output power.

Viewing Vital Product Data (VPD)

The VPD tab displays vital product data (if available) for the selected adapter port such as the product name, part number, serial number and so on.

To view VPD information:

1. Select **Host View** or **Fabric View**.
2. In the discovery tree, select the port whose VPD information you want to view.
3. Select the **VPD** tab.

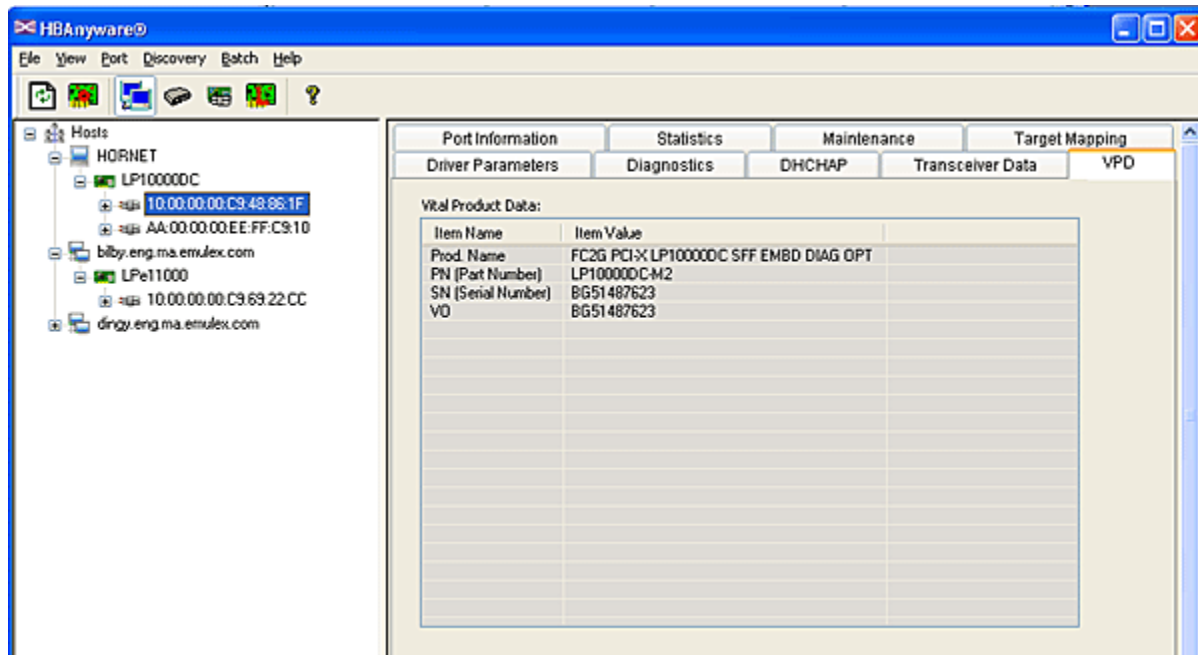


Figure 23: VPD tab

VPD Table Definitions

- Product Name - Displays product information about the selected adapter port.
- Part Number - Displays the adapter's part number.
- Serial Number - Displays the adapter's serial number.
- VO - Vendor unique data. "V" indicates a vendor-specific field. An adapter may have none, one or more of these fields defined. Valid values for this field are "VO" (the letter "O", not the number zero) and "Vx" (where "x" is a number).

Note: Solaris systems may show additional VPD information such as EC (EC level) and MN (manufacturer ID)

Viewing Maintenance Information

Use the Maintenance tab to view current firmware and WWPN and WWNN information for the selected adapter port. You can also update the adapter port's firmware, configure boot and change the port's WWPN and WWNN. (Not available in read-only mode.)

To view the firmware information:

1. Select **Host View** or **Fabric View**.
2. Select an adapter port in the discovery-tree.
3. Select the **Maintenance** tab.

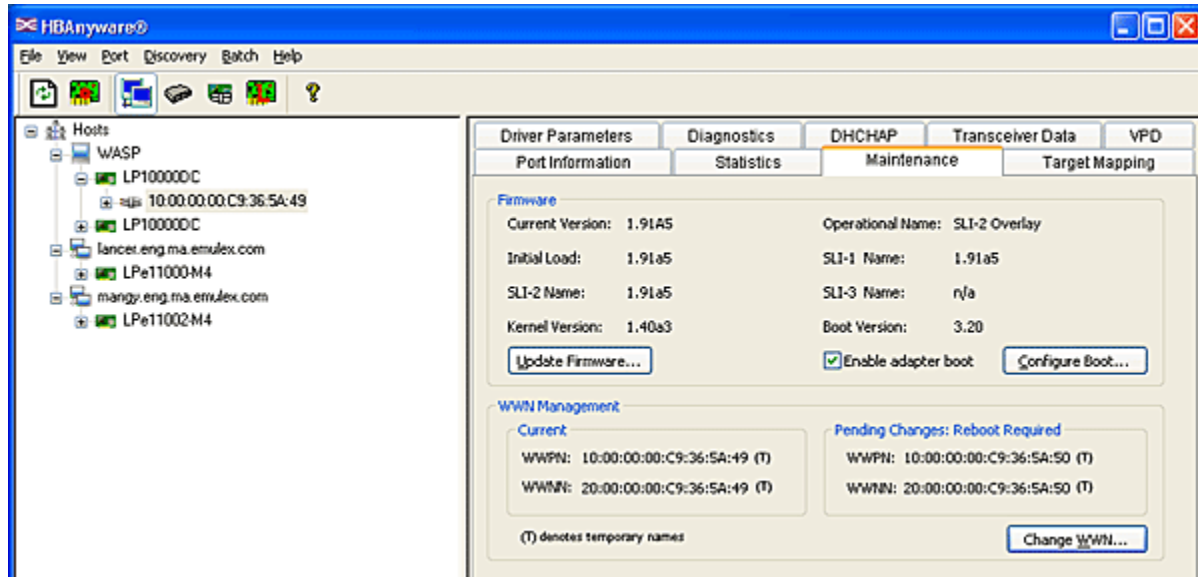


Figure 24: Maintenance Tab

Firmware Field Definitions

Firmware Area

- Current Version - The Emulex firmware version number for this model of adapter.
- Initial Load - The firmware version stub responsible for installing SLI code into its proper slot.
- SLI-2 Firmware Name - The name of the SLI-2 firmware overlay.
- Kernel Version - The version of the firmware responsible for starting the driver.
- Operational Name - The name of the operational firmware for the selected adapter.
- SLI-1 Firmware Name - The name of the SLI-1 firmware overlay.
- SLI-3 Firmware Name - The name of the SLI-3 firmware overlay.
- Boot Version - Displays one of the following:
 - The selected adapter port's boot code version if boot code is present.
 - "Disabled" if the boot code is disabled.
 - "Not Present" if boot code is not loaded. If boot code is not loaded, the Enable Adapter boot checkbox is not visible and you cannot configure the selected port to boot from SAN.

- Enable adapter boot checkbox - Check this box if you want the adapter to load and execute boot code during system startup. Click **Configure Boot** to configure boot from SAN. (Not available in read-only mode.) See “Configuring Boot from SAN” on page 106 for more information.

Note: Enabling adapter boot only causes the adapter to load the boot code and execute it during system startup. It does not mean that the adapter will boot from SAN. To boot from SAN, the boot type must be enabled to boot from SAN. Do this in the Boot from SAN configuration window for each boot type.

WWN Management Area

Current

- WWPN - Displays the World Wide Port Name for the selected adapter port.
- WWNN - Displays the World Wide Node Name for the selected adapter port.

Pending Changes

- WWPN - Works in conjunction with the Change WWN button. Displays the World Wide Port Name you assigned for the selected adapter port, but the system must be rebooted for these changes to take effect and appear under the “Current” listing. See “Changing World Wide Name Configuration” on page 144 for more information.
- WWNN - Works in conjunction with the Change WWN button. Displays the World Wide Node Name you assigned for the selected adapter port, but the system must be rebooted for these changes to take effect and appear under the “Current” listing. See “Changing World Wide Name Configuration” on page 144 for more information.

Firmware Tab Buttons (Not available in read-only mode.)

- Update Firmware - Click to update firmware on the selected adapter. See “Updating Firmware” on page 116 for more information.
- Configure Boot - Check Enable adapter boot and click Configure Boot to configure boot from SAN. See “Configuring Boot from SAN” on page 106 for more information.
- Change WWN - Click to change the selected adapter port's World Wide Node Name or World Wide Port Name. See “Changing World Wide Name Configuration” on page 144 for more information.

Viewing Target Information

Target Information contains information specific to the selected storage device.

To view target information:

1. Select **Host View**, **Fabric View** or **Virtual Port View**.
2. In the discovery-tree, select the target whose information you want to view. The Target Information tab appears.

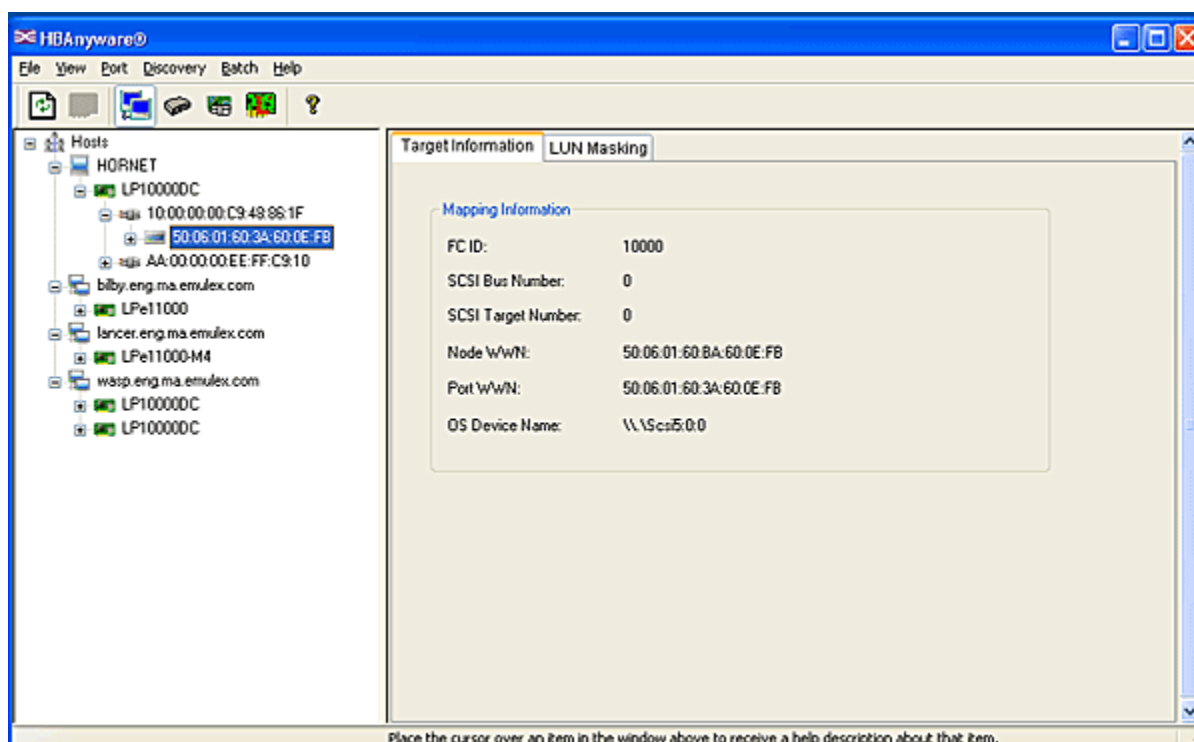


Figure 25: Target Information tab

Target Information Field Definitions

- Mapping Information Area
 - FC ID - The FC ID for the target; assigned automatically in the firmware.
 - SCSI Bus Number - The SCSI Bus number to which the target is mapped.
 - SCSI Target Number - The target's identifier on the SCSI Bus.
 - Node WWN - A unique 64-bit number, in hexadecimal, for the target (N_PORT or NL_PORT).
 - Port WWN - A unique 64-bit number, in hexadecimal, for the fabric (F_PORT or Switched Fabric Loop Port [FL_PORT]).
 - OS Device Name - The operating system device name.

Note: See "Masking and Unmasking LUNs (Windows, Solaris LPFC and Solaris SFS)" on page 130 for more information on LUN Masking.

Viewing LUN Information

The LUN Information tab contains information about the selected Logical Unit Number (LUN).

To view the LUN information:

1. Select **Host View**, **Fabric View** or **Virtual Port View**.
2. In the discovery-tree, select the LUN whose information you want to view. The LUN Information tab appears.

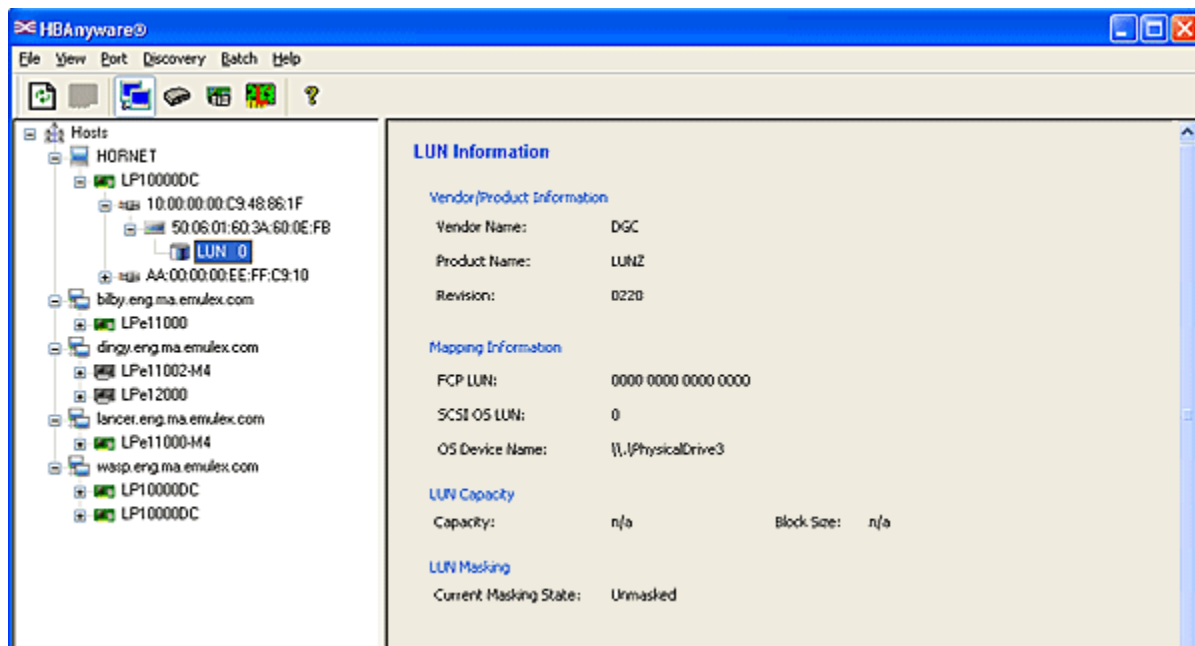


Figure 26: LUN Information

LUN Information Field Definitions

Vendor Product Information Area

- Vendor Name - The name of the vendor of the LUN.
- Product Name - The vendor-specific ID for the LUN.
- Revision - The vendor-specific revision number for the LUN.

Mapping Information Area

- FCP LUN - The FC identifier used by the adapter to map to the SCSI OS LUN.
- SCSI OS LUN - The SCSI identifier used by the OS to map to the specific LUN.
- OS Device Name - The name assigned by the OS to the LUN.

LUN Capacity Area

Note: LUN capacity information is only provided when the LUN is a mass-storage (disk) device. Other devices like tapes and scanners, etc. do not display capacity.

- Capacity - The capacity of the LUN, in megabytes.
- Block Size - The length of a logical unit block in bytes.

LUN Masking Area

- Current Masking State - Possible states are masked or unmasked.

Viewing Target Mapping (Windows, Solaris LPFC and Solaris SFS)

The Target Mapping tab enables you to view current target mapping and to set up persistent binding.

To view target mapping:

1. Select **Host View** or **Fabric View**.
2. In the discovery-tree, select the adapter port whose target mapping information you want to view.
3. Select the **Target Mapping** tab.

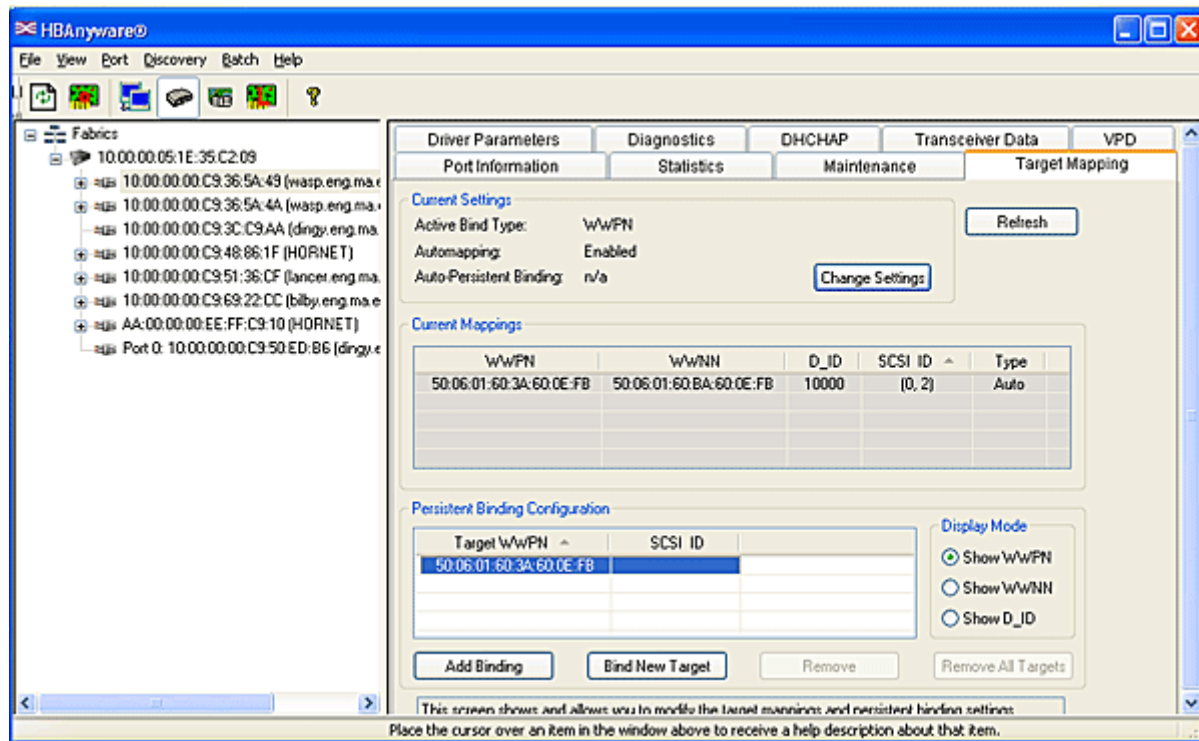


Figure 27: Target Mapping tab

Target Mapping Field Definitions

Current Settings Area

- Active Bind Type - WWPN, WWNN, or a destination identifier (D_ID).
- Automapping - The current state of SCSI device automapping: enabled (default) or disabled.
- Auto-Persistent Binding - The current state of the Auto-Persistent binding service.

Current Mappings Area

- This table lists current mapping information for the selected adapter port.

Persistent Binding Configuration Area

- This table lists persistent binding information for the selected adapter port.

Display Mode Radio Buttons

- Show WWPN, Show WWNN or Show DID options enable you to choose how to display information in the Persistent Binding Configuration table.

Target Mapping Buttons

- Refresh - Click to refresh the Target Mapping tab.
- Change Settings - Click to enable or disable automapping, choose a bind type and enable or disable LUN mapping and unmasking.
- Add Binding - Click to add a persistent binding.
- Bind New - Click to add a target that does not appear in the Persistent Binding table.
- Remove - Click to remove the selected binding.
- Remove All - Click to remove all persistent bindings that are displayed.

Viewing Target Mapping (Linux)

Use this tab to view target mapping. The Target Mapping tab is read-only.

Note: Persistent binding is not supported by the Linux 2.6 kernel or by the Emulex 8.0 or 8.2 versions of the driver for Linux.

To view target mapping:

1. Select **Host View** or **Fabric View**.
2. Select the adapter port in the discovery-tree whose target mapping information you want to view.
3. Select the **Target Mapping** tab.

Target Mapping Field Definitions

Current Settings Area

- Active Bind Type - N/A
- Automapping - N/A

Current Mappings Area

- This table lists current mapping information for the selected adapter.

Persistent Binding Configuration Area

- N/A

Display Mode Radio Buttons

- N/A

Target Mapping Buttons

- N/A

Creating and Deleting Virtual Ports

The Virtual Ports tab enables you to create and delete virtual ports.

Note: Creation and deletion of virtual ports is not supported on HBAs installed in VMware ESX Server machines.

Creating Virtual Ports


You can have the HBAnyware utility automatically generate the WWPN for the virtual port based on the WWPN for the physical port or you can manually type the WWPN. You can generate virtual ports on 4 Gb/s and 8 Gb/s HBAs. You cannot generate virtual ports on 1 Gb/s and 2 Gb/s HBAs.

The NPIV driver parameter must be enabled before attempting to create a virtual port. The driver parameter name varies slightly depending upon your operating system:

- For Windows: enableNPIV. On the Storport Miniport system, the SLIMode driver parameter must also be set to 0 or 3.
- For Solaris: enable-npiv
- For Linux 8.2: lpfc_enable_npiv

See “Configuring the Driver” on page 50 for more information on enabling driver parameters.

To create a virtual port:

1. Do one of the following:
 - From the **View** menu, select **Virtual Ports**.
 - From the toolbar, click  **Virtual Ports View**.
2. From the discovery-tree, select the adapter port on which you want to create a virtual port. The Virtual Ports tab appears.
3. Do one of the following:
 - Check **Auto-generate world wide port name**. The HBAnyware utility creates the unique WWPN for the new virtual port based on the WWPN of the physical port. This option allows you to automatically create up to 255 unique virtual ports for each physical port. It also has the advantage that the new WWPN is unique.

Note: After auto-generating 255 unique virtual ports, you cannot auto-generate anymore virtual ports even if you delete existing auto-generated ports. However, you can still enter your own World-Wide Port Name to create a virtual port.”

- Check **Use the following world-wide port name** and enter a unique WWPN you want to use. A valid port name must have one of the following formats:

```
10:00:xx:xx:xx:xx:xx:xx
2x:xx:xx:xx:xx:xx:xx:xx
3x:xx:xx:xx:xx:xx:xx:xx
5x:xx:xx:xx:xx:xx:xx:xx
```

where x is a hexadecimal value.

Caution: Ensure that a manually entered WWPN is unique to your particular SAN. Failure to do so could result in a non-functioning SAN and data loss.

4. Enter an optional name for the virtual port if you want. You can give the new virtual port any name you want up to 99 characters in length. This name is used as part of the Symbolic Node Name for the VPort.
5. Click **Create Virtual Port**. A dialog box appears notifying you that the virtual port was created. The dialog box also displays the new virtual port's WWPN. Each virtual port has its own WWPN, but its WWNN is the same as the physical port's WWNN.

Note: If you entered a WWPN that is already in use, you are prompted to enter another WWPN.

6. Click **OK**. The new virtual port is added to the discovery-tree under the physical port where it was created and the Number of Virtual Ports field is updated.

Note: The HBAnyware utility automatically refreshes its discovery after a virtual port is created. However, targets for a new virtual port may not be discovered during the refresh. Therefore, you must refresh the discovery until the targets appear under the virtual port in the discovery-tree.

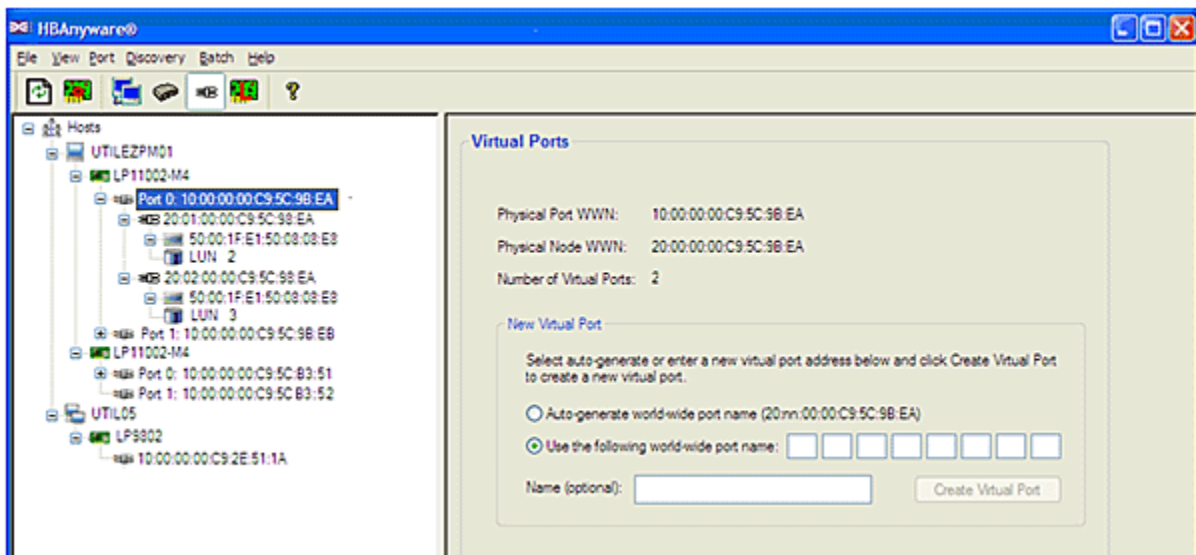


Figure 28: Virtual Ports window

Once you create a virtual port, a confirmation message similar to the following is displayed:

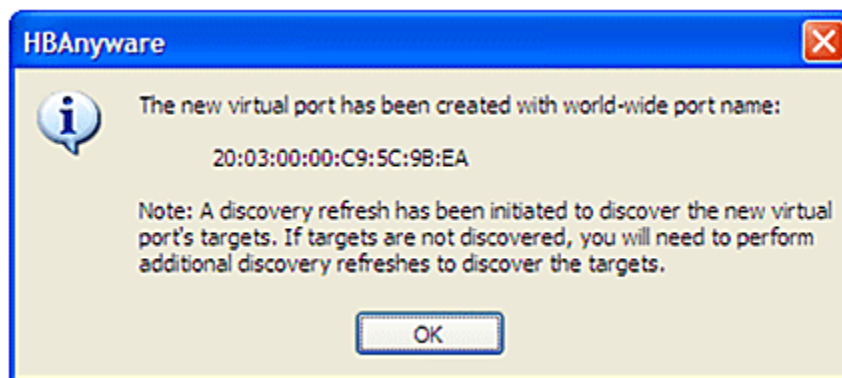



Figure 29: Successful VPort creation message

7. Click **OK**.

Deleting Virtual Ports

To delete a virtual port:

- Do one of the following:
 - From the **View** menu, select **Virtual Ports**.
 - From the toolbar, click  **Virtual Ports View**.
- From the discovery-tree, select the virtual port you want to delete. The Virtual Ports tab appears.

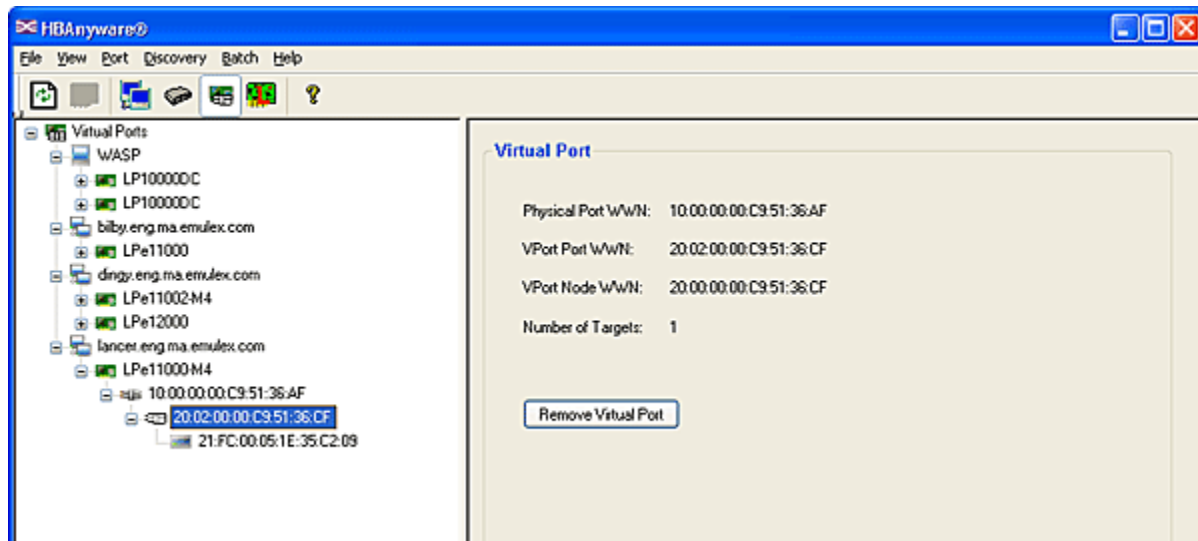


Figure 30: Virtual Port window

- Click **Remove Virtual Port**. The Delete Virtual Port Warning dialog box appears.

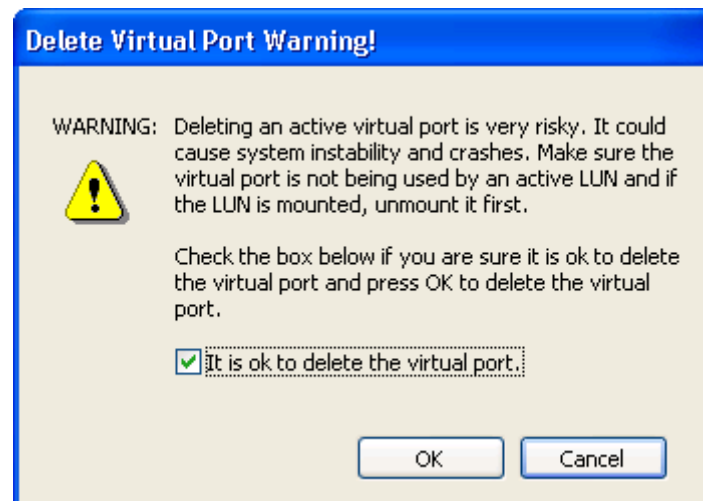


Figure 31: Delete Virtual Port Warning

Note: The link on the physical port must be up to delete a virtual port. The Remove Virtual button on the Virtual Port window is disabled if the link is down.

- Check **It is OK to delete the virtual port** and click **OK**. You are notified that the virtual port is no longer available and that it was removed from the discovery-tree.
- Click **OK**.

Configuring the Driver

In Windows, Solaris LPFC, Solaris SFS and Linux: Set driver parameters using the HBAnyware utility. In Solaris LPFC, Solaris SFS and Linux, you can also specify parameters when loading the driver manually. (Not available in read-only mode.) Refer to the appropriate driver manual for instructions.

Setting Driver Parameters

The Driver Parameters tab for adapters and hosts enable you to modify driver parameters for a specific adapter or all adapters in a host.

For example, if you select a host in the discovery-tree, you can globally change the parameters for all HBAs in that host. If you select an adapter port in the discovery-tree, you can change the `lpfc_use_adisc`, `lpfc_log_verbose` and the `lpfc_nodev_tmo` parameters for only that adapter.

For each parameter, the Driver Parameters tabs show the current value, the range of acceptable values, the default value, and the activation requirement. You can also restore parameters to their default settings.

Note: The global default driver parameters for converged network adapters (LP 21000 series CNAs) cannot be changed.

You can apply driver parameters for one adapter to other adapters in the system using the Driver Parameters tab, thereby simplifying multiple adapter configuration. See “Creating a Batch Mode Driver Parameters File” on page 58 for more information.

Note: The Linux 2.6 kernel only supports setting the `log_verbose`, `nodev_tmo` and `use_adisk` driver parameters for individual HBAs. You must apply other driver parameters to all HBAs contained in the host.

Note: For all compatible Linux versions: If you change driver parameters using the HBAnyware utility and you want these changes to be permanent and persist across system reboots, you must create a new ramdisk image. The ramdisk image is used when the kernel is initialized during system startup, and loads the LPFC driver with the updated driver parameters.

To create a new ramdisk you can use the LPFC driver's `lpfc-install` script. Refer to the “Creating a New Ramdisk” section of the Emulex Driver for Linux User Manual for instructions.

To change the driver parameters for a single adapter:

1. Select **Host View** or **Fabric View**.
2. In the discovery-tree, select the adapter port whose parameters you wish to change.

3. Select the **Driver Parameters** tab. The parameter values for the selected adapter are displayed.

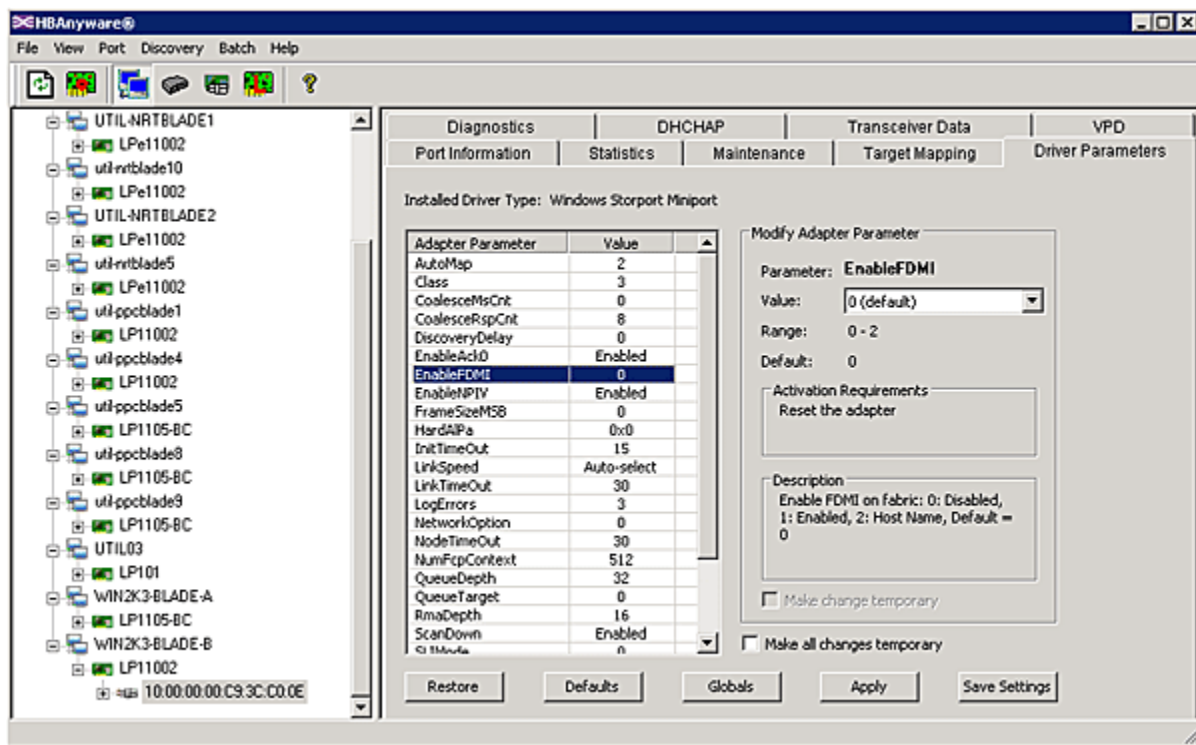


Figure 32: Driver Parameters tab - Adapter Selected

4. In the Driver Parameters tab, click the parameter that you want to change. A description of the parameter appears on the right side of the tab.
5. Enter a new value in the Value field in the same hexadecimal or decimal format as the current value or select a value from the drop-down menu. If you enter a value and the current value is in hexadecimal format, it is prefaced by "0x" (for example, 0x2d). You can enter a new hexadecimal value without the "0x". For example, if you enter ff10, this value is interpreted and displayed as "0xff10".
6. If you want the change to be temporary (causing the parameter to revert to its last permanent setting when the system is rebooted), check the **Make change temporary** box. This option is available only for dynamic parameters.
7. If you are making changes to multiple parameters, and you want all the changes to be temporary, check the **Make all changes temporary** box. This setting overrides the setting of the **Make change temporary** box. Only dynamic parameters can be made temporary.
8. Click **Apply** (or **Save**).

Restoring All Parameters to Their Earlier Values

If you changed parameters, but did not click **Apply** (or **Save**) and you want to restore the parameters to their last saved values, click **Restore**.

Resetting All Default Values

To reset all parameter values to their default (factory) values, click **Defaults**.

Setting an Adapter Parameter Value to the Host Adapter Parameter Value

Note: The global default driver parameters for converged network adapters (LP 21000 series CNAs) cannot be changed.

To set an adapter parameter value to the corresponding host parameter value:

1. Select **Host View** or **Fabric View**.
2. In the discovery-tree, select the adapter port.
3. Select the **Driver Parameters** tab.
4. Click **Globals**. All parameter values are now the same as the global, or host, values.
5. To apply the global values, click **Apply** (or **Save**).

Saving Adapter Driver Parameters to a File

To save adapter driver parameters, click **Save** (or **Save Settings**). Each definition is saved in a comma-delimited file with the following format:

```
<parameter-name>=<parameter-value>
```


The file is saved in the Emulex Repository directory. The HBAnyware utility can then use the Batch Driver Parameter Update function to apply these saved settings to any or all compatible HBAs on the SAN.

Note: Persistent binding settings cannot be saved with the Save (or Save Settings) feature.

Note: Host driver parameters cannot be saved.

Setting Driver Parameters for All HBAs in a Host

To change the driver parameters for all HBAs installed in a host:

1. Do one of the following:
 - From the **View** menu, click **Hosts**.
 - From the toolbar, click  **Host View**.
2. In the discovery-tree, click the host whose adapter driver parameters you want to change.
3. Select the **Driver Parameters** tab. If there are adapters with different driver types installed, the **Installed Driver Types** menu shows a list of all driver types and driver versions that are installed. Select the driver whose parameters you want to change. This menu does not appear if all the adapters are using the same driver.
4. In the Driver Parameters tab, click the parameter that you want to change. A description of the parameter appears on the right side of the tab.

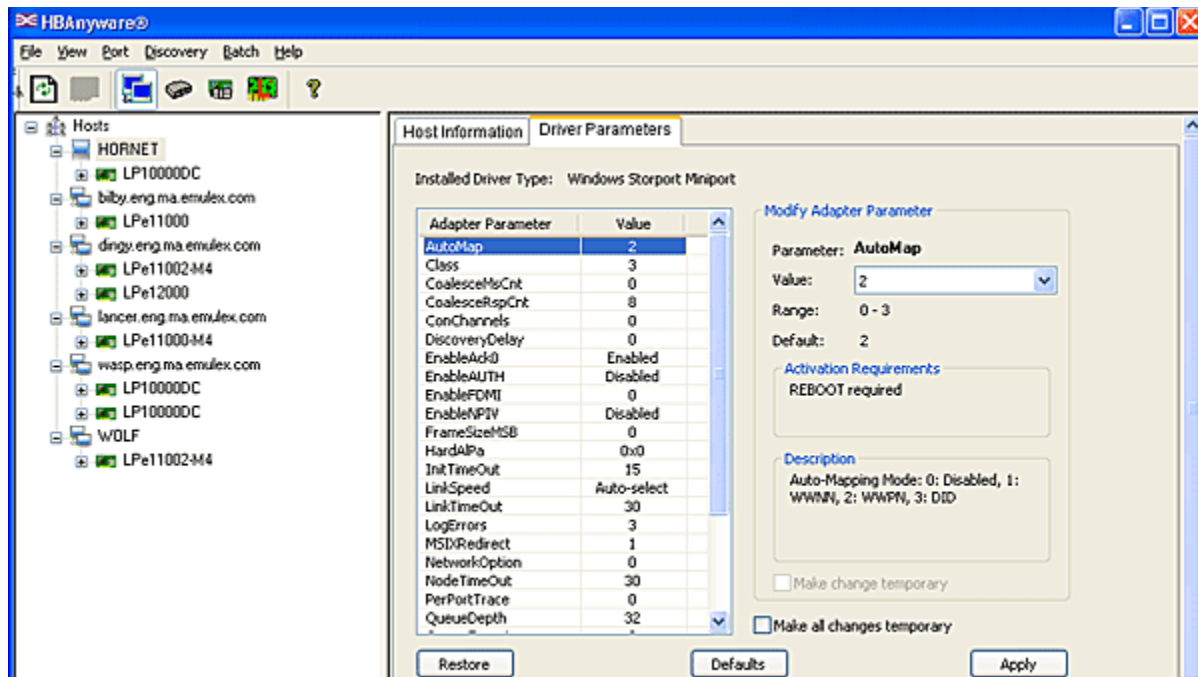


Figure 33: Driver Parameters tab - Host Selected

5. Enter a new value in the Value field in decimal or hexadecimal format, depending on how the current value is presented. If the value is in hexadecimal format, it is prefaced by "0x" (for example "0x2d").
6. To make a change temporary (the parameter to revert to its last permanent setting when the system is rebooted), check **Make changes temporary**. This option is available only for dynamic parameters.
7. To make changes to multiple parameters, check **Make all changes temporary**. Only dynamic parameters can be made temporary.
8. Click **Apply** (or **Save**).

Changing Non-dynamic Parameter Values (Linux)

To change non-dynamic parameter values for Linux version 8.0:

1. Navigate to the /usr/sbin/hbanyware directory and run the scripts to stop the HBAAnyware utility processes. Type:


```
./stop_hbanyware
```
2. Stop all I/O to LPFC attached devices.
3. Unload the lpfcdfc driver. Type:


```
modprobe -r lpfcdfc
```
4. Unload the LPFC driver. Type:


```
modprobe -r lpfc
```
5. Reload the driver. Type:


```
modprobe lpfc
modprobe lpfcdfc
```
6. Start the elxhbamgr service (remote service). Type:


```
./start_elxhbamgr
```

The HBAnyware utility discovery service starts automatically when you launch the application.

7. If the machine has the HBAnyware utility with Web Launch installed, the RMI services must be restarted. Type:

```
./start_weblaunch
```

To change non-dynamic parameter values for Linux version 8.2:

1. Navigate to the /usr/sbin/hbanyware directory and run the scripts to stop the HBAnyware utility processes. Type:

```
./stop_hbanyware
```

2. Stop all I/O to LPFC attached devices.

3. Unload the LPFC driver. Type:

```
modprobe -r lpfc
```

4. If DHCHAP authentication is currently employed on this machine, start up the Emulex Fibre Channel authentication service. Type:

5. /etc/init.d/fcauthd start

6. Reload the driver. Type:

```
modprobe lpfc
```

7. Start the elxhbamgr service (remote service). Type:

```
./start_elxhbamgr
```

The HBAnyware utility discovery service starts automatically when you launch the application.

Note: If DHCHAP authentication is currently employed on Emulex adapters on this machine, you must type "/etc/init.d/fcauthd start" to restart the authentication daemon.

8. If the machine has the HBAnyware utility with Web Launch installed, the RMI services must be restarted. Type:

```
./start_weblaunch
```

Note: For both Linux 8.0 and Linux 8.2, in order for changes to persist after a reboot, you must create a new ramdisk image.

Creating a Batch Mode Driver Parameters File

You can apply driver parameters for one adapter to other adapters in the system using the Driver Parameters tab. When you define parameters for an adapter, you create a .dpv file. The .dpv file contains parameters for that adapter. After you create the .dpv file, the HBAnyware utility enables you to assign the .dpv file parameters to multiple adapters in the system. (Not available in read-only mode.)

To create and assign the .dpv file:

1. Select **Host View** or **Fabric View**.
2. Select the adapter port whose parameters you want to apply to other adapters from the discovery-tree.
3. Select the **Driver Parameters** tab.
4. Set the driver parameters.
5. After you define the parameters for the selected adapter, click **Apply**.

6. Click **Save** (or **Save Settings**). The Select Driver Parameter File browse window appears.

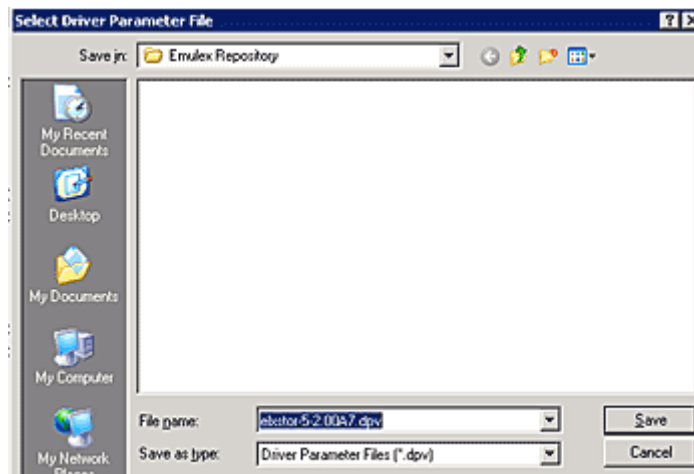


Figure 34: Select Driver Parameter File browse window

7. Use the Select Driver Parameter File dialog box to browse to where you want to save the file or to rename the file.
8. Click **Save**. The Save Driver Parameters dialog box appears.

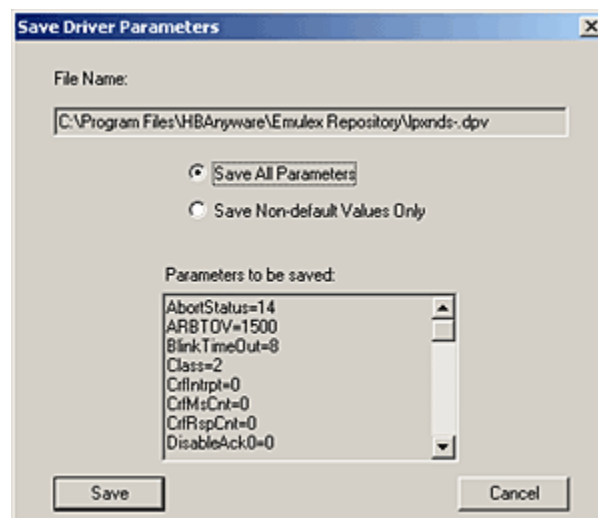


Figure 35: Save Driver Parameters dialog box

9. Use the two radio buttons to choose the type of parameters to save. You can save all parameters or only those parameters whose current values differ from their corresponding default values.

A list of the saved parameters and their current values show in the Saved Parameters box.

10. Click **Save**.

Assigning Batch Mode Parameters

To assign batch mode parameters to adapters:

1. From the **Batch** menu, select **Update Driver Parameters**. (You do not need to select any discovery-tree elements at this time.) The Select Driver Parameter File dialog box appears.
2. In Windows: Select the file whose parameters you want to apply and click **Open**.

In Solaris LPFC, Solaris SFS and Linux systems: A Browse button is included on the Batch Driver Parameters Update dialog box. The Browse button allows you to navigate to a different file.

Solaris systems: In addition to the Browse button, the Start Update, Reset Display, Save Log File and Close buttons are displayed.

The Batch Driver Parameter Update dialog box shows all the batch file compatible adapters with a check mark beside them.

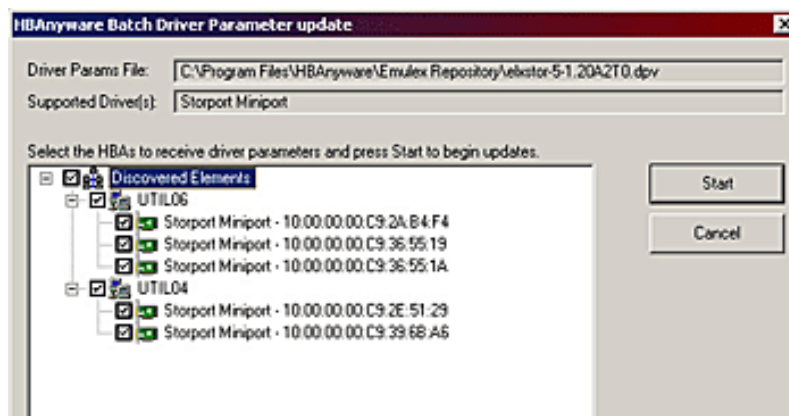


Figure 36: Batch Driver Parameters Update dialog box for Windows

3. Click **Start** (or **Start Update**). The HBAware utility Batch Driver Update dialog box shows the current status of the update. When the update completes, a final summary shows the number of adapters that were successfully processed, and the number of adapters for which one or more parameter updates failed.
4. If you want, click **Print Log** (or **Save Log File**) to print a report of the update.

Note: Printing is not supported in Linux.

Storport Miniport Driver Parameters

The parameter values listed in Table 3 are applicable to Storport Miniport driver versions 2.00 or later. If you are using a version previous to 2.00, see the Storport Miniport Driver User Manual for that version's parameter information.

Activation Requirements

A parameter has one of the following activation requirements:

- **Dynamic** - The change takes effect while the system is running.
- **Reset** - Requires an adapter reset from the utility before the change takes effect.
- **Reboot** - Requires reboot of the entire machine before the change takes effect. In this case, you are prompted to perform a reboot when you exit the utility.

The Driver Parameter table provides information such as the allowable range of values and factory defaults. Parameters can be entered in decimal or hexadecimal format.

Note: If you are creating custom unattended installation scripts, any driver parameter can be modified and included in the script.

Most parameters default to a setting that optimizes adapter performance.

Table 3: Storport Miniport Driver Parameters

Parameter	Definitions	Activation Requirement
AutoMap= n	<p>AutoMap controls the way targets are assigned SCSI IDs. Discovered targets are assigned persistent SCSI IDs according to the selected binding method. Persistent bindings do not take effect with the driver in stand-alone mode.</p> <p>If set to 0 = automap is disabled. Uses the HBAware utility to persistently set the SCSI address of a discovered FCP capable FC node (target). If set to 1 = automap by WWNN. If set to 2 = automap by WWPN. If set to 3 = automap by DID</p> <p>Value: 0 - 3 Default = 2</p>	Reboot
Class= n	<p>Class selects the class of service on FCP commands. If set to 2, class = 2. If set to 3, class = 3.</p> <p>Value: 2 - 3 Default = 3</p>	Dynamic
CoalesceMsCnt= n	<p>CoalesceMsCn specifies wait time in milliseconds to generate an interrupt response if CoalesceRspCnt has not been satisfied. Zero specifies an immediate interrupt response notification. A non-zero value enables response coalescing at the specified interval in milliseconds.</p> <p>Value: 0 - 63 (decimal) or 0x0 - 0x3F (hex) Default = 0 (0x0)</p>	Reset

Table 3: Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement
CoalesceRspCnt= n	CoalesceRspCn specifies the number of response entries that trigger an Interrupt response. Value: 0 - 255 (decimal) or 0x1 - 0xFF (hex) Default = 8 (0x8)	Reset
DiscoveryDelay= n	DiscoveryDelay controls whether the driver waits for 'n' seconds to start port discovery after link up. If set to 0 = immediate discovery after link up. If set to 1 or 2 = the number of seconds to wait after link-up before starting port discovery. Value: 0 - 2 seconds (decimal) Default = 0.	Dynamic
EnableAck0= n	Set to 1 to force sequence rather than frame level acknowledgement for class 2 traffic over an exchange. This applies to FCP data exchanges on IREAD and IWRITE commands. Value: 0 - 1 (decimal) Default = 0	Reset
EnableAUTH	EnableAUTH enables fabric authentication. This feature requires the authentication to be supported by the fabric. Authentication is enabled when this value is set to 1. Value: 0 - 1 Default = 0	Reboot
EnableFDMI= n	If set to 1, enables management server login on fabric discovery. This allows Fabric-Device Management Interface (FDMI) to operate on switches that have FDMI-capable firmware. If set to 2, FDMI operates and uses the host name feature of FDMI. Value: 0 -2 (decimal) Default = 0	Reset
EnableNPIV= n	If set to 1, enables N_Port_ID virtualization (NPIV). Requires NPIV supported firmware for the adapter. Value: 0 -1 Default = 0 (disabled) Note: To run the driver using NPIV or SLI-3 optimization, the firmware must be version 2.72a0 or later. If an earlier version is used, the driver runs in SLI-2 mode and does not support NPIV. Note: NPIV is not available on 1 Gb/s and 2 Gb/s HBAs.	Reset

Table 3: Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement
FrameSizeMSB= n	<p>FrameSizeMSB controls the upper byte of receive FrameSize if issued in PLOGI. This allows the FrameSize to be constrained on 256-byte increments from 256 (1) to 2048 (8).</p> <p>Value: 0 - 8 Default = 0</p>	Reset
HardALPA=0x n	<p>HardALPA allows the adapter to use a hard assigned loop address.</p> <p>Value: 0x00 - 0xEF (hex) Default = 0x00 (use soft addressing, or flash stored hard address value)</p> <p>Note: Only valid AL_PAs can be used.</p>	Reset
InitTimeout= n	<p>Determines the number of time-out seconds during driver initialization for the link to come up. If the link fails to come up by InitTimeout, driver initialization exits but is still successful. If the link comes up before InitTimeout, the driver sets double the amount for discovery to complete.</p> <p>Value: 5 -30 seconds or 0x5 - 0x1E (hex) Default = 15 seconds (0xF)</p>	Reboot
LinkSpeed= n	<p>LinkSpeed has significance only if the adapter supports speeds other than 1 Gb/s.</p> <p>Value: Auto-select, 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s Default = Auto-select</p> <p>Note: Setting this option incorrectly can cause the adapter to fail to initialize.</p>	Reset
LinkTimeOut= n	<p>LinkTimeOut applies to a private loop only. A timer is started on all mapped targets using the link timeout value. If the timer expires before discovery is re-resolved, commands issued to timed out devices returns a SELECTION_TIMEOUT. The Storport driver is notified of a Bus change event which leads to the removal of all LUNs on the timed out devices.</p> <p>Value: 1 - 500 seconds or 0x0 - 0xFE (hex) Default = 30 (0x1E)</p>	Dynamic

Table 3: Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement
LogErrors= n	<p>LogErrors determine the minimum severity level required to enable entry of a logged error into the system event log. Errors are classified as severe, malfunction or command level. A severe error requires user intervention to correct a firmware or adapter problem. An invalid link speed selection is an example of a severe error. A malfunction error indicates that the system has problems, but user intervention is not required. An invalid fabric command type is an example of a malfunction error. A command level error: an object allocation failure is an example of a command error.</p> <p>If set to 0, all errors are logged. If set to 1, command level errors are logged. If set to 2, malfunction errors are logged. If set to 3, severe errors are logged.</p> <p>Value: 0 - 3 Default = 3</p>	Dynamic
NodeTimeout= n	<p>The node timer starts when a node (i.e. discovered target or initiator) becomes unavailable. If the node fails to become available before the NodeTimeout interval expires, the OS is notified so that any associated devices (if the node is a target) can be removed. If the node becomes available before NodeTimeout expires the timer is canceled and no notification is made.</p> <p>Value: 1 - 255 seconds or 0x0 - 0xFF (hex) Default = 30 (0x1E)</p>	Dynamic
QueueDepth= n	<p>QueueDepth requests per LUN/target (see QueueTarget parameter). If you expect the number of outstanding I/Os per device to exceed 32, then you must increase to a value greater than the number of expected I/Os per device (up to a value of 254). If the QueueDepth value is set too low, a performance degradation can occur due to driver throttling of its device queue.</p> <p>Value: 1 - 254 or 0x1 - 0xFE (hex) Default = 32 (0x20)</p>	Dynamic
QueueTarget= n	<p>QueueTarget controls I/O depth limiting on a per target or per LUN basis.</p> <p>If set to 0 = depth limitation is applied to individual LUNs. If set to 1 = depth limitation is applied across the entire target.</p> <p>Value: 0 -1 or 0x0 - 0x1 (hex) Default = 0 (0x0)</p>	Dynamic

Table 3: Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement
RmaDepth= n	<p>RmaDepth sets the remote management buffer queue depth. The greater the depth, the more concurrent management controls can be handled by the local node.</p> <p>Value: 8 - 64, or 0x8 - 0x40 (hex) Default = 16 (0x10)</p> <p>Note: The RmaDepth driver parameter pertains to the functionality of the HBAnyware utility.</p>	Reboot
ScanDown= n	<p>If set to 0 = lowest AL_PA = lowest physical disk (ascending AL_PA order). If set to 1 = highest AL_PA = lowest physical disk (ascending SEL_ID order).</p> <p>Value: 0 - 1 Default = 0</p> <p>Note: This option applies to private loop only in DID mode.</p>	Reboot
SLIMode= n	<p>If set to 0 = autoselect firmware, use the newest firmware installed. If set to 2 = implies running the adapter firmware in SLI-2 mode. If set to 3 = implies running the adapter firmware in SLI-3 mode.</p> <p>Value: 0, 2 and 3 Default = 0</p>	Reboot
Topology= n	<p>Topology values can be 0 to 3. If set to 0 (0x0) = FC Arbitrated Loop (FC-AL). If set to 1 (0x1) = PT-PT fabric. If set to 2 (0x2) = *FC-AL first, then attempt PT-PT. If set to 3 (0x3) = *PT-PT fabric first, then attempt FC-AL.</p> <p>* Topology fail-over requires v3.20 firmware or higher. If firmware does not support topology fail-over, options 0,2 and 1,3 are analogous.</p> <p>Value: 0 - 3 Default = 2 (0x2)</p>	Reset
TraceBufSiz= n	<p>TraceBufSiz sets the size in bytes for the internal driver trace buffer. The internal driver trace buffer acts as an internal log of the driver's activity.</p> <p>Value: 250,000 - 2,000,000 or 0x3D090 - 0x1E8480 (hex). Default = 250,000 (0x3D090)</p>	Reboot

Table 4: Storport Miniport Topology Reference

Topology	Description	Value
Private Loop Operation	<p>Only FC-AL topology is used. After successful loop initialization, the driver attempts login with FL_PORT.</p> <ul style="list-style-type: none"> • If FL_PORT login is successful, public loop operation is employed. • If FL_PORT login is unsuccessful, private loop mode is entered. If a fabric is not discovered and the topology is arbitrated loop, the driver operates in private loop mode using the following rules: <ul style="list-style-type: none"> • If an FC-AL device map is present, each node described in the map is logged and verified as a target. • If an FC-AL device map is not present, logins are attempted with all 126 possible FC-AL addresses. LPGA/PRLO are also handled by the driver. Reception of either causes a new discovery or login to take place. 	0
Switched Fabric Operation	<p>Only switched F_PORT (point-to-point [pt.-to-pt.]) login is successful, fabric mode is used.</p> <ul style="list-style-type: none"> • If F_PORT login is unsuccessful, N_PORT-to-N_PORT direct connection topology will be used. • If a switch is discovered, the driver performs the following tasks: <ul style="list-style-type: none"> • FL_PORT login (Topology = 0;). • F_PORT login (Topology =1;). • Simple Name Server login. • State Change Registration. • Symbolic Name Registration. • FCP Type Registration if RegFcpType is set to 1.T • The driver logs out and logs in again. The name server indicates that registration is complete. • Simple Name Server Query for devices - the registry parameter SnsAll determines whether all N_Ports are requested (SnsALL=1;) or only SCSI FCP N_Ports (SnsAll=0; default) • Discovery/device creation occurs for each target device described by the Name Server. • The driver handles RSCN and LOGO/PRLO. Reception of either causes new discovery/logins to take place. 	1
*FC-AL attempt first, then attempt pt.-to-pt.	<ul style="list-style-type: none"> • Topology fail-over requires v3.20 firmware or higher. If firmware does not support topology fail-over, options 0 and 2 are analogous. Options 1 and 3 are analogous. 	2
*pt.-to-pt. fabric attempt first, then attempt FC-AL.	<ul style="list-style-type: none"> • Topology fail-over requires v3.20 firmware or higher. If firmware does not support topology, fail-over options 0 and 2 are analogous. Options 1 and 3 are analogous. 	3

Server Performance (Windows)

I/O Coalescing

I/O Coalescing is enabled and controlled by two driver parameters: CoalesceMsCnt and CoalesceRspCnt. The effect of I/O Coalescing depends on the CPU resources available on the server. With I/O Coalescing turned on, interrupts are batched, reducing the number of interrupts and maximizing the number of commands processed with each interrupt. For heavily loaded systems, this provides better throughput.

With I/O Coalescing turned off (the default), each I/O processes immediately, one CPU interrupt per I/O. For systems not heavily loaded, the default provides better throughput. The following table shows recommendations based upon the number of I/Os per adapter.

Table 5: Recommended Settings for I/O Coalescing

I/Os per Second	Suggested CoalesceMsCnt	Suggested CoalesceRspCnt
I/Os < 10000	0	8
10000 < I/Os < 18000	1	8
18000 < I/Os < 26000	1	16
I/Os > 26000	1	24

CoalesceMsCnt

The CoalesceMsCnt parameter controls the maximum elapsed time in milliseconds that the adapter waits before it generates a CPU interrupt. The value range is 0 - 63 (decimal) or 0x0 - 0x3F (hex). The default is 0 and disables I/O Coalescing.

CoalesceRspCnt

The CoalesceRspCnt parameter controls the maximum number of responses to batch before an interrupt is generated. If CoalesceRspCnt expires, an interrupt is generated for all responses collected up to that point. With CoalesceRspCnt set to less than 2, response coalescing is disabled and an interrupt is triggered for each response. The value range for CoalesceRspCnt is 1 - 255 (decimal) or 0x1 - 0xFF (hex). The default value is 8.

Note: A system reset is required to make changes to CoalesceMsCnt and/or CoalesceRspCnt.

Performance Testing

Four driver parameters must be considered (and perhaps changed from the default) for better performance testing: QueueDepth, NumFcpContext, CoalesceMsCnt and CoalesceRspCnt.

Note: Parameter values recommended in this topic are for performance testing only and not for general operation.

QueueDepth

If the number of outstanding I/Os per device is expected to exceed 32, increase this parameter to a value greater than the number of expected I/Os per device, up to a maximum of 254. The QueueDepth parameter defaults to 32. If set to a value that is not high enough, performance degradation may occur due to Storport throttling its device queue.

NumFcpContext

If the number of outstanding I/Os per adapter is expected to exceed 512, increase this parameter to a value greater than the number of expected I/Os per adapter. Increase this value in stages: from 128 to 256 to 512 to 1024 to a maximum of 2048. NumFcpContext limits the number of outstanding I/Os per adapter, regardless of how QueueDepth is set. The NumFcpContext defaults to 512. If NumFcpContext is too small relative to the total number of outstanding I/Os on all devices combined, performance degradation can occur due to I/O stream throttling.

CoalesceMsCnt

CoalesceMsCnt defaults to zero. If you are using a performance evaluation tool such as Iometer and if you expect the I/O activity to be greater than 8000 I/Os per second, set CoalesceMsCnt to 1 and re-initialize with an adapter reset or system reboot.

CoalesceRspCnt

CoalesceRspCnt defaults to 8. For all other values up to the maximum of 63, the adapter does not interrupt the host with a completion until either CoalesceMsCnt milliseconds has elapsed or CoalesceRspCnt responses are pending. The value of these two driver parameters reduces the number of interrupts per second which improves overall CPU utilization. However, there is a point where the number of I/Os per second is small relative to CoalesceMsCnt and this slows down the completion process, causing performance degradation.

Performance Testing Examples

Test Scenario One

You execute Iometer with an I/O depth of 1 I/O per device in a small-scale configuration (16 devices). In this case, the test does not exceed the adapter's performance limits and the number of I/Os per second are in the low thousands.

Recommendation: set CoalesceMsCnt to 0 (or leave the default value).

Test Scenario Two

You execute Iometer with an I/O depth of 48 per device in a small-scale configuration (16 devices).

Recommendation: set QueueDepth to be greater than 48 (e.g. 64) and NumFcpContext to be greater than 512 (e.g. 1024).

Driver for Solaris LPFC – Configuration File Reference

The parameter values listed in Table 6 are applicable to driver version 6.21g or later. If you are using a version previous to 6.21g, see the Emulex Driver for Solaris User Manual for parameter information.

Note: The fcp-bind-WWNN, fcp-bind-WWPN and fcp-bind-DID driver properties do not apply to a specific adapter. They are the global properties. These properties specify a list of persistent bindings. Each entry in this list applies to a specific instance of an adapter. You can only use one type of binding per adapter.

The LPFC.conf file contains all the driver properties that control driver initialization. In the LPFC.conf file, all adapter-specific driver properties have lpfcX-prefix (where X is the driver instance number); e.g., setting lpfc0-lun-queue-depth=20 makes 20 the default number of maximum commands which can be sent to a single logical unit (disk). The LPFC man page also provides further device property details.

Note: To override a driver parameter for a single driver-loading session, specify it as a driver property to the modload command. For example: # modload /kernel/drv/lpfc automap=0 (for 32-bit platforms) or modload /kernel/drv/sparcv9/lpfc automap=0 (for 64-bit platforms). This will load the Emulex SCSI support driver with automap set to 0 for this session.

Table 6: LPFC.conf Parameters

Property Name	Scope	Default	Min	Max	Dynamic	Comments
ack0	Controller Specific	0	0=Off	1=On	No	Use ACK0 for class 2. If ack0 is 1, the adapter tries to use ACK0 when running Class 2 traffic to a device. If the device doesn't support ACK0, then the adapter uses ACK1. If ack0 is 0, only ACK1 is used when running Class2 traffic.

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
auth-cfgparms	Controller Specific	<p>Description and Values: This is the DHCHAP related driver property for FC-SP support. It is only valid when driver property enable-auth is set to 1. This driver property is ignored when enable-auth is set to 0. The format of this property is:</p> <p>"LWWN RWWN auth_tov auth_mode bidir typelist hashlist dhgplist reauth_intval"</p> <p>LWWN: The WWPN of the local entity, i.e. adapter port. Use the form of NNNNNNNNNNNNNNNNN, where NNNNNNNNNNNNNNNNN is a 16 digit representation of the Host port World Wide Port Name. Or use 0000000000000000 to refer to local port WWPN.</p> <p>RWWN: The WWPN of the remote entity, i.e. fabric controller or any remote nport. Use the form of NNNNNNNNNNNNNNNNN, where NNNNNNNNNNNNNNNNN is a 16 digit representation of the fabric controller or nport FFFFFFFFFFFFFFFF as generic remote fabric port WWPN.</p> <p>auth_tov: The authentication timeout value in seconds. The atov range is 20 to 999 seconds in hexadecimal. For example, enter 45 seconds as 002d.</p> <p>auth_mod: The authentication mode. The valid modes are specified as 01 (Disabled), 02 (Enabled) and 03 (Passive). For detailed description of the mode, please refer to the Emulex HBAnyware (4.0) utility help page.</p> <p>bidir: The bi-directional authentication parameter. When set to 01, bidirectional authentication is enabled. When set to 00, bi-directional authentication is disabled. When bidirectional authentication is enabled, the key associated with remote entity must be specified in driver property authkeys.</p> <p>typelist: The authentication type list. Currently the Emulex lpfc driver only supports DHCHAP, tlist should always be set to 01000000.</p> <p>hashlist: The authentication hash list. Currently the Emulex lpfc driver only supports MD5 and SHA1. 01 refers to MD5, 02 refers to SHA1. For example: 01020000 means MD5, SHA1 in order of preference. 01000000 means MD5 only.</p> <p>dhgplist: The DHCHAP group list in order of preference. Currently the Emulex lpfc driver supports NULL DHCHAP algorithm and non-NULL DHCHAP algorithm such as DH group 1024, group 1280, group 1536 and group 2048. For example: 0102030405000000 means NULL, group 1024,1280, 1536 and 2048 in order of preference.</p> <p>reauth_intval: Reauthentication heartbeat interval in minutes. For example, 0000012c means the host side will do the reauthentication every 300 minutes. When set to 00000000 then reauthentication heartbeat is disabled.</p> <p>You can use lpfcX-auth-cfgparms to specify the per adapter instance DHCHAP authentication parameters setup. Any valid setup in this way will overwrite the auth-cfgparms setup.</p> <p>Note: If you are using a version previous to 6.21g, see the Emulex Driver for Solaris User Manual for the auth-cfgparms format.</p>				

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
auth-keys	Controller Specific	<p>Description and Values: This is the DHCHAP authentication key driver property for FC-SP support. It is only valid when driver property enable-auth is set to 1. This driver property should be ignored when enable-auth is set to 0. The format of this property is:</p> <p>"LWWN:type:length:pwd:RWWN:type:length:pwd"</p> <p>LWWN: The WWPN of the local entity, i.e. adapter port. You should use the form of NNNNNNNNNNNNNNNNN, where NNNNNNNNNNNNNNNNN is a 16 digit representation of the Host port WorldWide Port Name. Or you could use 0000000000000000 to refer to local port WWPN.</p> <p>type: The type of the key. The valid type could be ASCII text format represented by 0001, or binary format (Hexadecimal input) represented by 0002, or 0003 ignored. When 0003 is used, the corresponding klength and key are ignored. The format is 4 digits.</p> <p>length: The length of the key in bits. The length is represented by hexadecimal format. For example: 32 bytes of key should be represented by 0100. The maximum size of key is 128 bytes. The minimum size of key is 16 bytes.</p> <p>RWWN: The WWPN of the remote entity, i.e. Fabric controller or any remote port. You should use the form of NNNNNNNNNNNNNNNNN, where NNNNNNNNNNNNNNNNN is a 16 digit representation of the Fabric Controller or port FFFFFFFFFFFFFFFF as generic remote fabric port WWPN.</p> <p>pard: The key associated with local entity or remote entity. For example, 16 bytes of key with ASCII type: aabbccddeeffgghh. 16 bytes of key with binary type:61616262636364646565666667676868.</p> <p>You can use lpfcX-auth-keys to specify the per adapter instance DHCHAP authentication keys. Any valid setup in this way will overwrite the auth-keys setup.</p> <p>Note: If you are using a version previous to 6.21g, see the Emulex Driver for Solaris User Manual for auth-keys parameter format.</p>				

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
automap	Controller Specific	1	0=Off	1=On	No	Automatically assign SCSI IDs to FCP targets detected. If automap is 1, SCSI IDs for all FCP nodes without persistent bindings are automatically generated based on the bind method of the corresponding adapter port. If FCP devices are added to or removed from the FC network when the system is down, there is no guarantee that these SCSI IDs remain the same when the system is booted again. If automap is 0, only devices with persistent bindings are recognized by the system.
cr-count	Controller Specific	1	1	255	No	This value specifies a count of I/O completions after which an interrupt response is generated. This feature is disabled if cr-delay is set to 0.
cr-delay	Controller Specific	0	0	63	No	This value specifies a count of milliseconds after which an interrupt response generated if the cr-count has not been satisfied. This value is set to 0 to disable the Coalesce Response feature as default.
delay-rsp-err	Controller Specific	0	0=Off	1=On	Yes	(Boolean) The driver delays FCP RSP errors from being returned to the upper SCSI layer based on the no-device-delay configuration driver property.
discovery-threads	Controller Specific	1	1	32	No	Number of ELS commands during discovery. This value specifies the number of threads permissible during device discovery. A value of 1 serializes the discovery process.

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
dqfull-throttle-up-inc	Controller Specific	1	0	128	Yes	Amount to increment LUN queue depth each time. This driver property causes the LPFC driver to decrement a LUN's queue depth, if a queue-full condition is received from the target. The queue depth is decremented to a minimum of 1. The variables dqfull-throttle-up-inc and dqfull-throttle-up-time are used to restore the queue depth to the original value. The dqfull-throttle-up-time driver property defines a time, in seconds, that is used to tell when to increase the current queue depth. If the current queue depth isn't equal to the lun-queue-depth, and the driver stop_send_io flag is equal to 0 for that device, increment the current queue depth by dqfull-throttle-up-inc (don't exceed the lun-queue-depth). So, if both driver properties are set to 1, then the driver increments the current queue depth once per second until it equals the lun-queue-depth. The only other way to restore the queue depth (besides rebooting) to the original LUN throttle is by running the command /opt/lpfc/resetqdepth X. This restores the LUN throttle of all LUNs for adapter X to the original value.

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
dqfull-throttle-up-time	Controller Specific	1	0	30	Yes	Time interval (seconds) to increment LUN queue depth. This driver property causes the LPFC driver to decrement a LUN's queue depth, if a queue full condition is received from the target. The queue depth is decremented down to a minimum of 1. The variables dqfull-throttle-up-inc and dqfull-throttle-up-time are used to restore the queue depth back to the original. The dqfull-throttle-up-time driver property defines a time, in seconds, that is used to tell when to increase the current queue depth. If the current queue depth isn't equal to the lun-queue-depth, and the driver stop_send_io flag is equal to 0 for that device, increment the current queue depth by dqfull-throttle-up-inc (don't exceed the lun-queue-depth). So, if both driver properties are set to 1, then the driver increments the current queue depth once per second until it hits the lun-queue-depth. The only other way to restore the queue depth (besides rebooting), back to the original LUN throttle, is by running the command <code>/opt/lpfc/resetqdepth X</code> . This restores the LUN throttle of all LUNs for adapter X to the original value.
enable-auth	Controller Specific	0	0	1	Yes	This driver property specifies if the DHCHAP is enabled. When set to 1, also set up two other driver properties such as auth-cfgparms and auth-keys. When set to 0, DHCHAP support is disabled and auth-cfgparms and auth-keys are ignored. Any per adapter instance setup, for example, <code>lpfcX-enable-auth=1, 0</code> overwrites the value set by enable-auth.

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
extra-io-tmo	Controller Specific	0	0	255	Yes	Extra timeout value, in seconds, to be applied to each FCP command sent. When connecting through a large fabric, certain devices can require a longer timeout value.
fcg-bind-DID	Global	Inactive	N/A	N/A	No	Setup persistent FCP bindings based on a target device's Port ID. This binding guarantees that target assignments are preserved between reboots. The format for a bind entry is "NNNNNN:lpfcXtY" where NNNNNN is a 6 digit representation of the targets Port ID, X is the driver instance number and Y is the target assignment. Multiple entries must be separated by a comma (,) with the last entry terminated with a semi-colon (;). A sample entry follows: fcp-bind DID="0000ef:lpfc0t0"; (all on one line.)
fcg-bind-method	Controller Specific	2	1	4	No	Specifies the method of binding to be used. This binding method is used for persistent binding and automapped binding. A value of 1 forces WWNN binding, value of 2 forces WWPN binding and value of 3 forces DID binding. A fcp-bind-method value of 4 causes target ID assignment in a private loop environment to be based on the ALPA array (hard addressed). If a binding method is not specified for a port, WWPN binding is used. Any persistent binding whose method does not match with the bind method of the port is ignored. A sample entry follows: lpfc0-fcp-bind-method=1; lpfc1-fcp-bind-method=2;

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
fcplib-WWNN	Global	Inactive	N/A	N/A	No	Setup persistent FCP bindings based on a target device's WWNN. This binding guarantees that target assignments are preserved between reboots. The format for a bind entry is "NNNNNNNNNNNNNNNNNN:lpfcXtY" where NNNNNNNNNNNNNNNNNN is a 16 digit representation of the targets WorldWide Node Name, X is the driver instance number and Y is the target assignment. Multiple entries must be separated by a comma (,) with the last entry terminated with a semi-colon (;). A sample entry follows: fcplib-WWNN="20000020370c396f:lpfc1t0", "20000020370c27f7:lpfc0t2";
fcplib-WWPN	Global	Inactive	N/A	N/A	No	Setup persistent FCP bindings based on a target device's WWPN. This binding guarantees that target assignments are preserved between reboots. The format for a bind entry is "NNNNNNNNNNNNNNNNNN:lpfcXtY" where NNNNNNNNNNNNNNNNNN is a 16 digit representation of the targets WorldWide Port Name, X is the driver instance number and Y is the target assignment. Multiple entries must be separated by a comma (,) with the last entry terminated with a semi-colon (;). A sample entry follows: fcplib-WWPN="21000020370cf8263:lpfc1t0";
fcplib-class	Controller Specific	3	2	3	Yes	The LPFC driver is capable of transmitting FCP data in Class2 or Class 3. The LPFC driver defaults to using Class 3 transmission.

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
fdmi-on	Global	0	0	2	No	This driver property controls the FDMI capability of the LPFC driver. If set to 0 (default), FDMI is disabled. A value of 1 enables FDMI without registration of host name port attribute, while a value of 2 enables FDMI with registration of host name port attribute.
ip-class	Controller Specific	3	2	3	Yes	FC is capable of transmitting IP data in Class2 or Class 3. The LPFC driver defaults to using Class 3 transmission.
link-speed	Controller Specific	Auto-detect	Auto-Detect, 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s		No	Sets link speed.
linkdown-tmo	Controller Specific	30	0	255	Yes	This variable controls how long the driver holds I/O (0 - 255 seconds) after the link becomes inaccessible. When this timer expires, all I/O waiting to be serviced is aborted. For instance, FCP commands are returned back to the target driver with a failure. The lower the value, the quicker the driver fails commands back to the upper levels. There is a trade-off here: small values risk retrying the commands when the link is bouncing; large values risk delaying the failover in a fault tolerant environment. linkdown-tmo works in conjunction with nodev-tmo. I/O fails when either of the two timers expires.
log-only	Controller Specific	1	0	1	Yes	When set to 1, log messages are only logged to syslog. When set to 0, log messages are also printed on the console.

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
log-verbose	Controller Specific	0x0	0x0	0xffff	Yes	(bit mask) When set to non-zero this variable causes LPFC to generate additional messages concerning the state of the driver and the I/O operations it carries out. These messages can go to the system log file, /var/adm/messages and/or the system console. See Error Messages for detailed information on the bit mask.
lpfcXtY-lun-throttle	Controller Specific	none	1	128	No	The maximum number of outstanding commands to permit for any logical unit on a specific target. This value overrides lun-queue-depth.
lpfcXtY-tgt-throttle	Controller Specific	none	1	10240	No	The maximum number of outstanding commands to permit for any target, including all LUNs on that target. This value overrides tgt-queue-depth.
lpfcXtYIZ-lun-mask	Controller Specific	none	0	1	Yes	The driver uses this value to determine whether or not to expose discovered LUNs to the OS. When set to 1, the discovered LUN is masked and not reported to the OS. When set to 0, the discovered LUN is reported to the OS.
lpfcX-lun-unmask	Controller Specific	none	0	1	Yes	The driver uses this value to determine whether to override the LUN masking or not. When set to 1, all LUNs on all targets on the specified LPFC instance are reported to the OS regardless of their respective lunmask settings. When set to 0 (default), the override is not in effect.
lpfcXtY-lun-unmask	Controller Specific	none	0	1	Yes	The driver uses this value to determine whether to override the LUN masking or not. When set to 1, all LUNS on a specified target are reported to the OS regardless of their respective lunmask settings. When set to 0 (default), the override is not in effect.

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
lun-queue-depth	Global	30	1	128	No	The driver uses this value as the default limit for the number of simultaneous commands to issue to a single logical unit on a single target on the loop. A single logical unit is never sent more commands than allowed by lun-queue-depth; however, less can be sent when sd-max-throttle or tgt-queue-depth is reached for the entire target.
msi-mode	Controller Specific	3	0	3	No	This variable controls whether LPFC uses MSI-based interrupts or legacy interrupts. If set to 3 (default), the driver tries to use multiple message MSI. If multiple message MSI is not possible due to an OS or hardware limitation, then the driver attempts single message MSI. If single message MSI fails, then the driver attempts legacy interrupts. A value of 0 disables MSI and the driver uses legacy interrupts.
network-on	Controller Specific	0	0	1	No	This variable controls whether LPFC provides IP networking functionality over FC. This variable is a Boolean: when zero, IP networking is disabled; when non-zero, IP networking is enabled.
no-device-delay	Global	1	0	30	Yes	This variable (0 to 30 seconds) determines the length of the interval between deciding to fail an I/O because there is no way to communicate with its particular device (e.g., due to device failure or device removal) and actually failing the command. A value of zero implies no delay whatsoever. This delay is specified in seconds. A minimum value of 1 (1 second) is recommended when communicating with any Tachyon based device.

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
nodev-holdio	Controller Specific	0	0=Off	1=On	Yes	This variable controls if I/O errors are held by the driver if a FCP device on the SAN disappears. If set, I/O errors are held until the device returns back to the SAN (potentially indefinitely). This driver property is ignored, if SCSI commands are issued in polled mode. The upper layer can retry the command once the error is returned.
nodev-tmo	Controller Specific	30	0	255	Yes	This variable controls how long I/O is held by the driver if a device on the SAN disappears. If set, I/O is held for the specified number of seconds. If the device does not appear on the SAN before nodev-tmo seconds, then the driver fails all held I/O and mark the device as unavailable. The upper layer can retry the command once the error is returned.
num-bufs	Controller Specific	128	64	4096	No	This variable specifies the number of command buffers to allocate. These buffers are used for FC Extended Link Services (ELS), and one for each FCP command issued in SLI-2 mode. If you want to queue lots of FCP commands to the adapter, then increase num-bufs for better performance. These buffers consume physical memory and are also used by the device driver to process loop initialization and re-discovery activities. Important: The driver must always be configured with at least several dozen ELS command buffers; Emulex recommends at least 128.

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
num-iocbs	Controller Specific	256	128	10240	No	This variable indicates the number of Input/Output Control Block (IOCB) buffers to allocate. IOCBs are internal data structures used to send and receive I/O requests to and from the LightPulse hardware. Too few IOCBs can temporarily prevent the driver from communicating with the adapter, thus lowering performance. (This condition is not fatal.) If you run heavy IP traffic, increase num-iocbs for better performance.
post-ip-buf	Controller Specific	128	64	1024	No	This variable specifies the number of 4K STREAMS buffers to allocate and post to the FC IP ring. Increase this setting for better IP performance under heavy loading.
scan-down	Controller Specific	1	0=Off	1=On	Yes	There are two scanning algorithms used to discover a node in a private loop. If scan-down is 1, devices on the private loop are scanned starting from ALPA 0x01 through ALPA 0xEF. If scan-down is 0, devices on the private loop are scanned starting from ALPA 0xEF through ALPA 0x01. Scan-down values 0 and 1 do not apply if a loop map is obtained. See the FC-AL profile for the definition of a loop map.
target-disk	Controller Specific	"sd"			No	Controls the FCP device to target driver associations. By default, FCP devices are associated with their corresponding Solaris native target drivers. FCP disks are associated with sd. To use a third party target driver, modify the corresponding entries. For example (all on one line): target-disk="scsidisk";

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
target-tape	Controller Specific	"st"			No	Controls the FCP device to target driver associations. By default, FCP devices are associated with their corresponding Solaris native target drivers. FCP tapes are associated with st. To use a third party target driver, modify the corresponding entries. For example: target-tape="IBMtape";
target-tapechanger	Controller Specific	"sgen"			No	Controls the FCP device to target driver associations. By default, FCP devices are associated with their corresponding Solaris native target drivers. FCP medium changers are associated with sgen. To use a third party target driver, modify the corresponding entries. For example: targettapechanger="IBMtape";
tgt-queue-depth	Global	0	0	10240	No	The driver uses this value as the default limit for the number of simultaneous commands to issue to a single target on the loop. A value of 0 causes no target throttling to occur. A single target is never be sent more commands than allowed by tgt-queue-depth; however, less can be sent when sd-max-throttle is reached for the entire target.

Table 6: LPFC.conf Parameters (Continued)

Property Name	Scope	Default	Min	Max	Dynamic	Comments
topology	Controller Specific	0x0	0x0=loop , then P2P 0x2=P2P only 0x4=loop only 0x6=P2P, then loop		No	This variable controls the FC topology expected by LPFC at boot time. FC offers point-to-point, fabric, and arbitrated loop topologies. To make the adapter operate as an N_Port, select point-to-point mode (used for N_Port to F_Port, and N_Port to N_Port connections). To make the adapter operate in a FC loop as an NL_Port, select loop mode (used for private loop and public loop topologies). The driver rejects an attempt to set the topology to a value not in the above list. The auto-topology settings 0 and 6 do not work unless the adapter is using firmware version 3.20 or higher.
use-adisc	Controller Specific	0	0=Off	1=On	Yes	This variable controls the ELS command used for address authentication during re-discovery upon link-up. If set, ADISC is used, otherwise, PLOGI is used. For FCP-2 devices, the driver always uses ADISC. For re-discovery due to a RSCN, the driver always uses ADISC.
xmt-que-size	Controller Specific	256	128	10,240	No	This variable specifies the number of network packets that can be queued or outstanding at any time in the driver. Increase this setting for better IP performance under heavy loading.

Driver For Solaris SFS Driver Parameters

- The emlxs.conf file contains all the parameters necessary to initialize the Solaris SFS driver.
- The HBAnyware utility reflects the Solaris SFS driver parameters.

The parameter values listed in Table 7 are applicable to driver versions 1.31/2.31 or later. If you are using the SFS driver version 1.22/2.22, see the HBAnyware Utility, version 3.3 User Manual. If you are using an earlier version of the SFS driver, see the Emulex Driver for Solaris User Manual for parameter information. All parameters are controller-specific.

Table 7: emlxs.conf Parameters

Property Name	Default	Min	Max	Activation	Comments
ack0	0	0	1	Dynamic	Use ACK0 for class 2. If ACK0 is 1, the adapter tries to use ACK0 when running Class 2 traffic to a device. If the device doesn't support ACK0, then the adapter uses ACK1. If ACK0 is 0, only ACK1 is used when running Class 2 traffic.
adisc-support	1	0	2	Dynamic	Sets the level of driver support for the FC ADISC login I/O recovery method. 1= Partial support. Flush I/O's for non-FCP2 target devices at link down 0 = No support. Flush active I/O's for all FCP target devices at link down. 2 = Full support. Hold active I/O's for all devices at link down.
assign-alpa	0x00	0x00	0xef	Link reset	This parameter is only valid if topology is set to loop. A 0x00 setting means no preference. If multiple adapter instances on the same host are on the same loop, set this value differently for each adapter.

Table 7: emlxs.conf Parameters (Continued)

Property Name	Default	Min	Max	Activation	Comments
auth-cfgs	<p>Description and Values: This is the DHCHAP related driver property for FC-SP support. It is only valid when driver property enable-auth is set to 1. This driver property is ignored when enable-auth is set to 0.</p> <p>This property represents a table of entries. The format of the table is:</p> <pre>"LWWN:RWWN:atov:amod:dir tlist:hlist:dhgplist:reauth", "LWWN:RWWN:atov:amod:dir tlist:hlist:dhgplist:reauth", "LWWN:RWWN:atov:amod:dir tlist:hlist:dhgplist:reauth";</pre> <p>The table can hold as many entries as needed.</p> <p>LWWN: The WWPN of the local entity, i.e. adapter port. You should use the form of NNNNNNNNNNNNNNNN, where NNNNNNNNNNNNNNNN is a 16 digit hexadecimal representation of the Host port World Wide Port Name. Or you could use 0000000000000000 to refer to the local port WWPN.</p> <p>RWWN: The WWPN of the remote entity, i.e. fabric controller. Use the form of NNNNNNNNNNNNNNNN, where NNNNNNNNNNNNNNNN is a 16 digit hexadecimal representation of the fabric controller or FFFFFFFFFFFFFFFF as generic remote fabric port WWPN.</p> <p>atov: The authentication timeout value in seconds (hexadecimal format). The atov range is 20 to 999 seconds in hexadecimal. For example, 45 seconds would be entered as 002d.</p> <p>amod: The authentication mode. The valid modes are specified as 1 (Disabled), 2 (Active) and 3 (Passive). For a detailed description of the mode, refer to the Emulex HBAnyware utility help page.</p> <p>dir: The bi-directional authentication parameter. When set to 1, bi-directional authentication is enabled. When set to 0, bi-directional authentication is disabled. When bidirectional authentication is enabled, the key associated with remote entity must be specified in the auth-keys driver property.</p> <p>tlist: The authentication type list (4 digits). Currently the Emulex LPFC driver only supports DHCHAP, tlist should always be set to 1000.</p> <p>hlist: The authentication hash list (4 digits). Currently the Emulex emlxs driver supports only MD5 and SHA1. 01 refers to MD5, 02 refers to SHA1. For example: 01020000 means MD5, SHA1 in order of preference. 01000000 means MD5 only.</p> <p>dhgplist: The DHCHAP group list in order of preference (8 digits). Currently Emulex emlxs driver supports NULL DHCHAP algorithm and non-NULL DHCHAP algorithm such as DH group 1024, group 1280, group 1536 and group 2048. The values can be 0 (undefined), 1 (NULL group), 2 (1024), 3 (1280), 4(1536), 5 (2048). For example: 12345000 means NULL, group 1024,1280, 1536 and 2048 in order of preference.</p> <p>reauth: Reauthentication heart beat interval in minutes (hexadecimal format). For example, 12c means the host side does the reauthentication every 300 minutes. When set to 0 reauthentication heartbeat is disabled. You can use emlxsX-auth-cfgs to specify the per adapter instance DHCHAP authentication parameters setup. Any valid setup in this way overwrites the auth-cfgs setup.</p>				

Table 7: emlxs.conf Parameters (Continued)

Property Name	Default	Min	Max	Activation	Comments
auth-cfgs (continued)	However, since the parameter represents a table of entries the table can represent all entries across all adapter instances. This allows all adapter instances to share a common table of entries.				
auth-keys	<p>Description and Values: This is the DHCHAP authentication key driver property for FC-SP support. It is only valid when driver property enable-auth is set to 1. This driver property should be ignored when enable-auth is set to 0. This property represents a table of entries. The format of the table is:</p> <pre>"LWWN:RWWN:Ltype:Lkey:Rtype:Rkey" , "LWWN:RWWN:Ltype:Lkey:Rtype:Rkey" , "LWWN:RWWN:Ltype:Lkey:Rtype:Rkey" ,</pre> <p>The table can hold as many entries as needed.</p> <p>LWWN: The WWPN of the local entity, that is an adapter port. You should use the form of NNNNNNNNNNNNNNNN, where NNNNNNNNNNNNNNNN is a 16 digit representation of the Host port WorldWide Port Name. Or you could use 0000000000000000 to refer to local port WWPN.</p> <p>RWWN: The WWPN of the remote entity, that is a fabric controller. Use the form of NNNNNNNNNNNNNNNN, where NNNNNNNNNNNNNNNN is a 16 digit representation of the fabric controller or FFFFFFFFFFFFFFFF as generic remote fabric port WWPN.</p> <p>Ltype: The local key type. The type field can be 1 (ASCII text formatted key) or 2 (binary hex formatted key).</p> <p>Lkey: The local key to be associated with the local entity. For example a key of ASCII type could look like: abcdefgh. A key of binary type could look like: 12ef58c98274d46.</p> <p>Rtype: The remote key type. The type field can be 1 (ASCII text formatted key) or 2 (binary hex formatted key).</p> <p>Rkey: The remote key to be associated with the remote entity. For example a key of ASCII type could look like: abcdefgh. A key of binary type could look like: 12ef58c98274d46.</p> <p>You can use emlxsX-auth-keys to specify the per adapter instance DHCHAP authentication keys. Any valid setup in this way overwrites the auth-keys setup. However, since the parameter represents a table of entries the table can represent all entries across all adapter instances. This allows all adapter instances to share a common table of entries.</p>				
console-notices	0x00000000	0x00000000	0xFFFFFFFF	Reboot	Verbose mask for notice messages to the console.
console-warnings	0x00000000	0x00000000	0xFFFFFFFF	Reboot	Verbose mask for warning messages to the console.
console-errors	0x00000000	0x00000000	0xFFFFFFFF	Reboot	Verbose mask for error messages to the console.

Table 7: emlxs.conf Parameters (Continued)

Property Name	Default	Min	Max	Activation	Comments
cr-count	1	1	255	Link reset	This value specifies a count of I/O completions after which an interrupt response is generated. This feature is disabled if cr-delay is set to 0.
cr-delay	0	0	63	Link reset	This value specifies a count of milliseconds after which an interrupt response generated if cr-count has not been satisfied. This value is set to 0 to disable the Coalesce Response feature as default.
enable-auth	0	0	1	Link reset	This driver property specifies if the DHCHAP is enabled or not. When set to 1, DHCHAP is enabled. When set to 0, DHCHAP support is disabled.
enable-npiv	0	0	1	Adapter reset	Enables NPIV support in the driver.
link-speed	Auto-Detect	Auto-Detect, 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s		Link reset	Sets link speed for initializing FC connection.
linkup-delay	10	0	60	Link reset	Sets the linkup delay period (seconds) after adapter initialization.
log-notices	0xFFFFFFFF	0x00000000	0xFFFFFFFF	Reboot	Verbose mask for notice messages to the messages file.
log-warnings	0xFFFFFFFF	0x00000000	0xFFFFFFFF	Reboot	Verbose mask for warning messages to the messages file.
log-errors	0xFFFFFFFF	0x00000000	0xFFFFFFFF	Reboot	Verbose mask for error messages to the messages file.

Table 7: emlxs.conf Parameters (Continued)

Property Name	Default	Min	Max	Activation	Comments
max-xfer-size	339968	131072	1388544	Reboot	<p>Sets the maximum SCSI transfer size in bytes per IO. This parameter is only used by the driver on i386 platforms. The driver does not limit transfer size on SPARC platforms. This parameter determines the scatter gather list buffer size. A pool of buffers is reallocated by the driver during boot. A larger transfer size requires a larger memory allocation.</p> <p>Memory_model/max-xfer-size Small/131072 - 339968 Medium/339969 - 688128 Large/688129 - 1388544</p> <p>Note: This parameter is supported in 2.31 version. (not 2.30 version).</p>
network-on	0	0	1	Reboot	Enables/disables IP networking support in the driver.
num-iocbs	1024	128	10240	Adapter reset	This variable indicates the number of Input/Output Control Block (IOCB) buffers to allocate.
num-nodes	0	0	4096	Adapter reset	Number of FC nodes (NPorts) the driver-supports.
pci-max-read	2048	512	4096	Reload the emlxs driver	Sets the PCI-X max memory read byte count [512, 1024, 2048 or 4096]
pm-support	0	0	1	Reboot	<p>Enable/Disable power management support in the driver.</p> <p>0 = Disables power management support in the driver.</p> <p>1 = Enables power management support in the driver.</p>
ub-bufs	1000	40	16320	Reboot	Sets the number of unsolicited buffers to be allocated.

Table 7: emlxs.conf Parameters (Continued)

Property Name	Default	Min	Max	Activation	Comments
target-mode	0	0	1	Reboot	Enables/Disables COMSTAR target mode support. If target mode is enabled for that port, then SFS initiator mode is disabled for that port
topology	0	0 =loop, then P2P 2 =P2P only 4 =loop only 6 =P2P, then loop		Link reset	Set to point-to-point mode if you want to run as an N_Port. Set to loop mode if you want to run as an NL_Port.
vport	Virtual port registration table. The enable-npiv must be set to 1. The vport table may have any number of comma delimited entries. Each entry must be of the form: "PHYS_WWPN:VPORT_WWNN:VPORT_WWPN:VPORT_ID" PHYS_WWPN = World Wide Port Name of adapter's physical port VPORT_WWNN = Desired World Wide Node Name of virtual port VPORT_WWPN = Desired World Wide Port Name of virtual port VPORT_ID = Desired virtual port id (1 to max vports) The port ids must start at 1 and increment by 1 with no gaps in the count. The virtual port id 0 is reserved for the physical port. Example: vport= "10000000c9123456:28010000c9123456:20010000c9123456:1", "10000000c9123456:28020000c9123456:20020000c9123456:2", "10000000c9123457:28010000c9123457:20010000c9123457:1", "10000000c9123457:28020000c9123457:20020000c9123457:2", "10000000c9123457:28030000c9123457:20030000c9123457:3"; All entries are automatically created or removed by the HBAnyware utility.				
vport-restrict-login	1	0	1	Reboot	Sets the virtual port's behavior when discovering targets in the SAN. 1 prevents the VPort from logging into other initiator ports on the SAN. Also rejects logins from other ports in the SAN because it assumes that all ports that send a PLOGI are Initiators. When this sysfs entry is turned off (0) the driver attempts to login to every port that it can access in the SAN and accept logins from all ports. NOTE: In a SAN where there are other initiators this feature greatly reduces the hardware resources the driver uses.

Solaris SFS and Solaris LPFC Driver Parameter Cross-Reference

The cross-reference information listed in Table 8 is applicable to the Driver for Solaris LPFC version 6.20i and the Driver for Solaris SFS version 1.30/2.30. If you are using Solaris SFS version 1.22.2.22 see the HBAnyware utility version 3.3 User manual. If you are using a Solaris LPFC or Solaris SFS driver version previous to these listed, see the appropriate user manual for parameter information.

Table 8: Solaris SFS and Solaris LPFC Driver Parameter Cross-Reference

Solaris emlxs/ HBAnyware Property	Solaris emlxs/HBAnyware Min/Max, Defaults and Description	Related LPFC Property	LPFC Min/Max, Default and Description	Comments
ack0	0 = Off 1 = On Default: 0 Description: Use ACK0 for class 2. If ACK0 is 1, the adapter tries to use ACK0 when running Class 2 traffic to a device. If the device doesn't support ACK0, then the adapter uses ACK1. If ACK0 is 0, only ACK1 is used when running Class 2 traffic.	ack0	0 = Off 1 = On Default: 0 Description: Use ACK0 for class 2. If ACK0 is 1, the adapter tries to use ACK0 when running Class 2 traffic to a device. If the device doesn't support ACK0, then the adapter uses ACK1. If ACK0 is 0, only ACK1 is used when running Class 2 traffic.	N/A
adisc-support	0 = No support. Flush active I/O's for all FCP target devices at link down. 1 = Partial support. Flush I/O's for non-FCP2 target devices at link down. 2 = Full support. Hold active I/O's for all devices at link down. Default: 1 Description: Sets the level of driver support for the FC ADISC login I/O recovery method.	use-adisc	0 = Off 1 = On Default: 0 Description: Controls the ELS command used for address authentication during rediscovery upon link-up. The driver always uses ADISC for FCP-2 devices and re-discovery due to an registered state change notification (RSCN).	If there are tape devices on the SAN that support FCP2, set the use-adisc property to 1 and the adisc-support property to 1 (partial support) or 2 (full support).
assign-alpa	Min:0x00 Max:0xef Default:0x00 (valid ALPA's only) Description: This is only valid if topology is loop. A zero setting means no preference. If multiple adapter instances on the same host are on the same loop, set this value differently for each adapter.	N/A	N/A	N/A

Table 8: Solaris SFS and Solaris LPFC Driver Parameter Cross-Reference (Continued)

Solaris emlxs/HBAnyware Property	Solaris emlxs/HBAnyware Min/Max, Defaults and Description	Related LPFC Property	LPFC Min/Max, Default and Description	Comments
auth-cfgs	This is the DHCHAP related driver property for FC-SP support. It is only valid when driver property enable-auth is set to 1. This driver property is ignored when enable-auth is set to 0.	auth-cfgparms	This is the DHCHAP related driver property for FC-SP support. It is only valid when driver property enable-auth is set to 1. This driver property is ignored when enable-auth is set to 0.	For detail on this emlxs parameter, see Table 7. For detail on this LPFC parameter, see Table 6.
auth-keys	This is the DHCHAP authentication key driver property for FC-SP support. It is only valid when driver property enable-auth is set to 1. This driver property is ignored when enable-auth is set to 0.	auth-keys	This is the DHCHAP authentication key driver property for FC-SP support. It is only valid when driver property enable-auth is set to 1.	For detail on this emlxs parameter, see Table 7. For detail on this LPFC parameter, see Table 6.
console-notices	Min: 0x00000000 Max: 0xFFFFFFFF Default: 0x00000000 Verbose mask for notice messages to the console.	log-verbose	Min: 0x0 Max: 0xffff Default: 0x0 (bit mask) When set to nonzero this variable causes LPFC to generate additional messages concerning the state of the driver and the I/O operations it carries out. These messages can go to the system console.	N/A
console-warnings	Min: 0x00000000 Max: 0xFFFFFFFF Default: 0x00000000 Verbose mask for warning messages to the console.	log-verbose	Min: 0x0 Max: 0xffff Default: 0x0 (bit mask) When set to nonzero this variable causes LPFC to generate additional messages concerning the state of the driver and the I/O operations it carries out. These messages can go to the system console.	N/A
console-errors	Min: 0x00000000 Max: 0xFFFFFFFF Default: 0x00000000 Verbose mask for error messages to the console.	log-verbose	Min: 0x0 Max: 0xffff Default: 0x0 (bit mask) When set to nonzero this variable causes LPFC to generate additional messages concerning the state of the driver and the I/O operations it carries out. These messages can go to the system console.	N/A

Table 8: Solaris SFS and Solaris LPFC Driver Parameter Cross-Reference (Continued)

Solaris emlxs/HBAnyware Property	Solaris emlxs/HBAnyware Min/Max, Defaults and Description	Related LPFC Property	LPFC Min/Max, Default and Description	Comments
cr-delay	Min:0 Max:63 Default:0 Description: Specifies a count of milliseconds after which an interrupt response is generated if the cr-count has not been satisfied. This value is set to 0 to disable the Coalesce Response feature as default.	cr-delay	Min:0 Max:63 Default:0 Description: Specifies a count of milliseconds after which an interrupt response is generated if the cr-count has not been satisfied. This value is set to 0 to disable the Coalesce Response feature as default.	Setting this value can minimize CPU utilization by reducing the number of interrupts that the driver generates to the operating system.
cr-count	Min:1 Max:255 Default:1 Description: Specifies a count of I/O completions after which an interrupt response is generated. This feature is disabled if cr-delay is set to 0.	cr-count	Min:1 Max:255 Default:1 Description: Specifies a count of I/O completions after which an interrupt response is generated. This feature is disabled if cr-delay is set to 0.	The value is often determined by your OEM. This property sets the number of I/Os to be queued in the operating system's driver before an interrupt is initiated. The driver default settings are roughly a 1:1 I/O to interrupt ratio. If you change this property, performance varies per application.
enable-auth	Min:0 Max:1 Default:0 This driver property specifies if the DHCHAP is enabled or not.	enable-auth	Min:0 Max:1 Default:0 This driver property specifies if the DHCHAP is enabled or not.	This parameter is dynamic for LPFC. This property requires a link reset for SFS.
link-speed	Auto-Detect, 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s Default: Auto-Detect Description: Sets the link speed setting for initializing the FC connection.	link-speed	Auto-Detect, 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s Default: Auto-Detect Description: Sets link speed.	This value can be changed to a specific link speed to optimize link initialization process for a specific environment.

Table 8: Solaris SFS and Solaris LPFC Driver Parameter Cross-Reference (Continued)

Solaris emlxs/HBAnyware Property	Solaris emlxs/HBAnyware Min/Max, Defaults and Description	Related LPFC Property	LPFC Min/Max, Default and Description	Comments
log-notices	Min: 0x00000000 Max: 0xFFFFFFFF Default: 0x00000000 Verbose mask for notice messages to the messages file.	log-verbose	Min: 0x0 Max: 0xffff Default: 0x0 (bit mask) When set to nonzero this variable causes LPFC to generate additional messages concerning the state of the driver and the I/O operations it carries out. These messages can go to the system log file, /var/adm/messages.	N/A
log-warnings	Min: 0x00000000 Max: 0xFFFFFFFF Default: 0x00000000 Verbose mask for warning messages to the messages file.	log-verbose	Min: 0x0 Max: 0xffff Default: 0x0 (bit mask) When set to nonzero this variable causes LPFC to generate additional messages concerning the state of the driver and the I/O operations it carries out. These messages can go to the system log file, /var/adm/messages.	N/A
log-errors	Min: 0x00000000 Max: 0xFFFFFFFF Default: 0x00000000 Verbose mask for error messages to the messages file.	log-verbose	Min: 0x0 Max: 0xffff Default: 0x0 (bit mask) When set to nonzero this variable causes LPFC to generate additional messages concerning the state of the driver and the I/O operations it carries out. These messages can go to the system log file, /var/adm/messages.	N/A
max-xfer-size	Min: 131072 Max: 1388544 Default: 339968 Determines the scatter gather list buffer size. A pool of buffers is reallocated by the driver during boot. A larger transfer size requires a larger memory allocation.	max-xfer-size	N/A	N/A

Table 8: Solaris SFS and Solaris LPFC Driver Parameter Cross-Reference (Continued)

Solaris emlxs/HBAnyware Property	Solaris emlxs/HBAnyware Min/Max, Defaults and Description	Related LPFC Property	LPFC Min/Max, Default and Description	Comments
network-on	Min:0 (Disables) Max:1 (Enables) Default:1 Description: Enables or disables IP networking support in the driver.	network-on	Min:0 (Disables) Max:1 (Enables) Default: 0 Description: Controls whether LPFC provides IP networking functionality over FC. This variable is Boolean: when zero, IP networking is disabled: when non-zero, IP networking is enabled.	The LPFC parameter enables or disables FCIP on the Emulex adapter.
num-iocbs	Min:128 Max:10240 Default = 1024 Description: Sets the number of iocb buffers to allocate.	num-iocbs	Min:128 Max:10240 Default = 256 Description: Specifies the This variable indicates the number of Input/Output control block (IOCB) buffers to allocate. IOCBs are internal data structures used to send and receive I/O requests to and from the LightPulse hardware. Too few IOCBs can temporarily prevent the driver from communicating with the adapter, thus lowering performance. (This condition is not fatal.) If you run heavy IP traffic, you should increase num-iocbs for better performance.	
num-nodes	Min:0 Max:4096 Default: 0 Description: Number of FC nodes (NPorts) the driver supports.	N/A	N/A	
pci-max-read	Min: 512 Max: 4092 Default: 2048 Description: Sets the PCI-X max memory read byte count [512, 1024, 2048 or 4096].	N/A	N/A	

Table 8: Solaris SFS and Solaris LPFC Driver Parameter Cross-Reference (Continued)

Solaris emlxs/HBAnyware Property	Solaris emlxs/HBAnyware Min/Max, Defaults and Description	Related LPFC Property	LPFC Min/Max, Default and Description	Comments
pm-support	0 = Disables power management support in the driver. 1 = Enables power management support in the driver. Default: 0 Description: Enable/Disable power management support in the driver.	N/A	N/A	
topology	0 = loop, if it fails attempt pt-to-pt 2 = pt-to-pt only 4 = loop only 6 = pt-to-pt, if it fails attempt loop Default: 0 Description: Link topology for initializing the FC connection. Set pt-to-pt if you want to run as an N_Port. Set loop if you want to run as an NL_Port.	topology	0x0 = loop, if it fails attempt pt-to-pt 0x2 = pt-to-pt only 0x4 = loop only 0x6 = pt-to-pt, if it fails attempt loop Default: 0 Description: Controls the FC topology expected by LPFC at boot time. FC offers pt-to-pt, fabric and arbitrated loop topologies. To make the adapter operate as an N_Port, select pt-to-pt mode (used for N_Port to F_Port and N_Port to N_Port connections). To make the adapter operate as an NL_Port, select loop mode (used for private loop and public loop topologies). The driver rejects an attempt to set the topology to a value not in the above list. The auto-topology settings 0 and 6 does not work unless the adapter is using firmware version 3.20 or higher.	The topology property controls protocol (not physical) topology attempted by the driver.
ub-bufs	Min:40 Max:16320 Default:1000 Description: Sets the number of unsolicited buffers to be allocated.	N/A	N/A	

Driver for Linux Parameter Tables

The driver parameter values listed in Table 9, Table 10 and Table 11 are applicable to driver version 8.0.16.34. If you are using a version previous to 8.0.16.34, see the *Emulex Driver for Linux User Manual* for parameter information.

Note: For all compatible Linux versions: If you change driver parameters using the HBAnyware utility and you want these changes to be permanent and persist across system reboots, you must create a new ramdisk image. The ramdisk image is used when the kernel is initialized during system startup, and loads the LPFC driver with the updated driver parameters.

To create a new ramdisk you can use the LPFC driver's lpfc-install script. Refer to the "Creating a New Ramdisk" section of the *Emulex Driver for Linux User Manual* for instructions.

The driver parameters listed in Table 12 and Table 13 are applicable to driver version 8.2. Driver for Linux version 8.2 supports DHCHAP authentication and configuration.

Version 8.0 LPFC and LPFCDFC Parameter

The parameters determine some aspects of the driver behavior. The following tables list the driver parameters. Some driver parameters can be modified and take effect only on a driver load while others can be modified dynamically and take effect immediately. The tables also list the default, minimum and maximum values for these parameters.

In Table 9, driver parameters marked with an asterisk are not supported by the HBAnyware utility. You can change them via LPFC. See the driver user manual for more information.

Table 9: Driver for Linux, LPFC Static Parameters (Requires a driver reload to change)

Variable	Default	Min	Max	Comments	Visible using sysfs
lpfc_ack0	0	0=Off	1=On	Use ACK0 for class 2.	Yes
lpfc_cr_count	1	1	255	This parameter determines the values for I/O coalescing for cr_delay (msec) or cr_count outstanding commands.	No
lpfc_cr_delay	0	0	63	This parameter determines the values for I/O coalescing for cr_delay (msec) or cr_count outstanding commands.	No

Table 9: Driver for Linux, LPFC Static Parameters (Requires a driver reload to change) (Continued)

Variable	Default	Min	Max	Comments	Visible using sysfs
lpfc_discovery_threads	32	1	64	<p>Specifies the maximum number of ELS commands that can be outstanding for a discovery.</p> <p>Note: The discovery_threads parameter defaults to a value of 64 for private loop topologies regardless of the configured value. If there are multiple ports configured on the host the value of 64 is only used for those ports that are connected in a private loop topology. The configured value is used for all other ports.</p> <p>Note: On the VMware ESX Server only, range is 30 - 64. The default is erroneously set to 1. Set this parameter to a valid value. If you do not, when you change any parameter value, the following error message can be displayed: "Driver Parameter 'discovery-threads' is not within the allowed range".</p>	No
lpfc_fcp_class	3	2	3	FC class for FCP data transmission.	Yes
lpfc_link_speed	0	0=auto select 1=1 Gb/s 2=2 Gb/s 4=4 Gb/s 8=8 Gb/s		Sets link speed.	Yes
lpfc_hba_queue_depth*	8192	32	8192	Maximum number of FCP commands that can queue to an Emulex adapter.	Yes
lpfc_lun_queue_depth	30	1	128	Default max commands sent to a single logical unit (disk).	Yes
lpfc_topology	0	0x0=loop then P2P 0x2=P2P only 0x4=loop only 0x6=P2P then loop		FC link topology (defaults to loop, if it fails attempts point-to-point mode).	Yes
lpfc_fcp_bind_method	2	1	4	<p>Specifies method of binding each port. Values:</p> <p>1: WWNN binding 2: WWPN binding 3: D_ID binding 4: ALPA binding</p>	Yes

Table 9: Driver for Linux, LPFC Static Parameters (Requires a driver reload to change) (Continued)

Variable	Default	Min	Max	Comments	Visible using sysfs
lpfc_fdm_i_on	0	0	2	False (0) if disabled. (1) or (2) if enabled depending on type of support needed.	Yes
lpfc_scan_down	1	0=Off	1=On	Select method for scanning ALPA to assign a SCSI ID.	Yes
lpfc_max_luns	256	1	32768	Specifies the maximum number of LUNs per target. A value of 20 means LUNs from 0 to 19 are valid.	Yes
lpfc_multi_ring_support*	1	1	2	Determines the number of primary SLI rings over which to spread IOCB entries.	No

* Variable not tunable in the HBAnyware utility.

All LPFC dynamic parameters are read/write using sysfs.

Table 10: Driver for Linux, LPFC Dynamic Parameters (Do not require a driver reload to change)

Variable	Default	Min	Max	Comments
lpfc_discovery_min_wait*	3	0	60	The minimum number of seconds the driver waits for the discovery to complete.
lpfc_discovery_wait_limit*	600	0	600 (special value meaning no limit)	The maximum number of seconds the driver waits for the discovery to complete.
lpfc_linkup_wait_limit*	15	0	60	The number of seconds the driver waits for the link to come up.
lpfc_log_verbose	0x0	0x0	0xffff	(bit mask) Extra activity logging.
lpfc_nodev_tmo	30	0	255	Seconds to hold I/O errors if device disappears.
lpfc_use_adisc	0	0=Off	1=On	Send ADISC instead of PLOGI for device discovery or RSCN.

Table 11: LPFCDFC Driver for Linux, Static Parameters

Variable	Default	Min	Max	Comments
lpfc_scsi_req_tmo	30	0	255	Time out value (in seconds) for SCSI request sent through lpfcdfc module. (Not available using the HBAnyware utility GUI. Command line only.)

Version 8.2 LPFC Parameters

DHCHAP Authentication and Configuration

The Emulex driver for Linux version 8.2.0.x supports the FC-SP/Authentication DHCHAP (Diffie-Hellmann Challenge Handshake Authentication Protocol). To activate FC-SP/Authentication between the adapter host port and fabric F_port using DHCHAP, you modify the DHCHAP associated driver properties in the driver configuration file.

The Emulex driver for Linux version 8.2.0.x supports MD5 and SHA-1 hash functions and supports the following DH groups: Null, 1024, 1280, 1536 and 2048.

Note: This version of the driver supports for N-Port to F-Port authentication only and does not support N-Port to N-Port authentication.

Enabling Authentication

Enabling authentication is a two step process. To enable authentication:

- The `fcauthd` daemon must be running.
- The `lpfc_enable_auth` module parameter must be set to enabled.

The `lpfc_enable_auth` Module Parameter

Use the `lpfc_enable_auth` module parameter to enable or disable authentication support. This module parameter can be set when loading the driver to enable or disable authentication on all Emulex adapters in the system, or it can be set dynamically after the driver is loaded to enable or disable authentication for each port (physical and virtual). The default setting for the `lpfc-enable-auth` module parameter is disabled. Refer to Table 13 starting on page 102 for the parameter values.

The `fcauthd` Daemon

The Emulex LPFC driver requires the `fcauthd` daemon to perform authentication tasks for it. To enable authentication you must have this daemon running. If you want to load the driver with authentication enabled, the `fcauthd` daemon should be running prior to driver load. The driver can start with authentication enabled if the daemon is not running, but all ports are placed into an error state. When the daemon is started the driver should discover the daemon and reset the adapter to enable the driver to perform authentication. To test if this daemon is running, start the daemon, or stop the daemon, you must use the `/etc/init.d/fcauthd` script. This script accepts the standard daemon parameters: start, stop, reload, status, restart and condrestart.

The script syntax is `/etc/init.d/fcauthd <parameter>`.

Note: The 8.2.0.X driver connects directly to the `fcauthd` daemon. To unload the driver you must first stop the `fcauthd` daemon. This will close the netlink connection and allow the LPFC driver to unload.

`fcauthd` Daemon Parameters

The `fcauthd` daemon supports the following parameters:

- `start` - To start the `fcauthd` daemon pass the start command to the `fcauthd` script. This command loads the daemon into memory, opens a netlink connection for the driver, and reads the authentication configuration database into memory for use by the LPFC driver.
- `stop` - To stop the `fcauthd` daemon pass the stop command to the `fcauthd` script. This command takes down the netlink connection between the `fcauthd` daemon and the `lpfc` driver, and stops the `fcauthd` daemon.

- reload - The reload command reloads the authentication configuration database into memory. This is done whenever the database is changed by another application (the HBAnyware utility) or by you. If the database is changed the new configuration information is not used until the fcauthd daemon reloads the database.
- status - This command is used to display the current status of the fcauthd daemon. The status should be either running or stopped.
- restart - The restart command performs a stop and then a start.
- condrestart - The conditional restart command checks the status of the fcauthd daemon. If it is running it issues a stop and then a start command. If the fcauthd daemon is not running nothing happens.

Setting Remote and Local Passwords

You can configure each port's password. See "Changing Your Password" on page 116 for more information.

Parameter Tables

The driver parameters determine some aspects of the driver behavior. The following tables list the driver parameters. Some driver parameters can be modified and take effect only on a driver load while others can be modified dynamically and take effect immediately. The tables also list the default, minimum and maximum values for these parameters.

Table 12: lpfc Static Parameters (Requires a driver reload to change)

Variable	Default	Min	Max	Comments	Visible using sysfs
lpfc_ack0	0	0=Off	1=On	Uses ACK0 for class 2.	Yes
lpfc_dev_loss_initiator	0	0	1	Engage devlos timeout for initiators	Yes
lpfc_discovery_threads	32	1	64	<p>Specifies the maximum number of ELS commands that can be outstanding for a discovery.</p> <p>NOTE: The discovery_threads parameter defaults to a value of 64 for private loop topologies regardless of the configured value. If there are multiple ports configured on the host the value of 64 is only used for those ports that are connected in a private loop topology. The configured value is used for all other ports.</p>	No

Table 12: lpfc Static Parameters (Requires a driver reload to change) (Continued)

Variable	Default	Min	Max	Comments	Visible using sysfs
lpfc_enable_da_id	0	0 = Disabled (default) 1 = enable – a DA_ID CT command will be sent to the fabric when logging out.		This parameter controls whether the driver will issue a DA_ID CT command to the fabric when vports logout of the fabric.	Yes
lpfc_enable_hba_heartbeat	1	0 = heartbeat disabled 1 = heartbeat enabled		Controls the adapter heartbeat logic in the driver. If the heartbeat is enabled and the heartbeat logic detects that the adapter is nonfunctional, the driver will shutdown the adapter.	No
lpfc_enable_hba_reset	1	0 = hba reset disabled 1 = hba reset enabled		Controls whether hba_resets will be allowed by the driver to pass to the adapter. This is used as a debugging tool.	No
lpfc_enable_npiv	0	0	1	This parameter controls the driver's ability to use NPIV to create virtual ports. It defaults to off (0) which prevents the driver from creating any virtual ports. When enabled (set to 1) it enables you to create and delete virtual ports (if supported by the fabric).	Yes
lpfc_fcp_class	3	2	3	The Fibre Channel class for FCP data transmission.	Yes
lpfc_hba_queue_depth	8192	32	8192	The maximum number of FCP commands that can queue to an Emulex adapter.	Yes
lpfc_lun_queue_depth	30	1	128	The default maximum commands sent to a single logical unit (disk).	Yes
lpfc_scan_down	1	0=Off	1=On	Selects method for scanning ALPA to assign a SCSI ID.	Yes
lpfc_sg_seg_cnt	64	64	256	Controls the max scatter gather segment count passed to the driver.	Yes. Displayed as sg_tablesize

Table 12: lpfc Static Parameters (Requires a driver reload to change) (Continued)

Variable	Default	Min	Max	Comments	Visible using sysfs
lpfc_sli_mode	0	0 = auto (default) 2 = SLI 2 mode 3 = SLI 3 mode (only available on newer HBAs)		This parameter allows you to force the SLI mode requested by the adapter driver.	No
lpfc_max_luns	256	1	32768	Specifies the maximum number of LUN IDs per target. A value of 20 means LUN IDs from 0 to 19 are valid. The SCSI layer scans each target until it reaches the specified LUN ID.	Yes
lpfc_multi_ring_rctl	4	1	255	Identifies RCTL for additional ring configuration. NOTE: Only used when multi_ring_support is enabled.	Yes
lpfc_multi_ring_support	1	1	2	Determines the number of primary SLI rings over which to spread IOCB entries.	No
lpfc_multi_ring_type	5	1	255	Identifies TYPE for additional ring configuration. NOTE: Only used when multi_ring_support is enabled.	Yes
lpfc_use_msi	0	0 = MSI disabled 1 = MSI enabled 2 = MSI-X enabled		Controls whether the driver uses Message Signaled Interrupts.	Yes

All lpfc dynamic parameters are read/write using sysfs.

Table 13: lpfc Dynamic Parameters (Do not require a driver reload to change)

Variable	Default	Min	Max	Comments
lpfc_cr_count	1	1	255	This parameter determines the values for I/O coalescing for cr_delay (msec) or cr_count outstanding commands.
lpfc_cr_delay	0	0	63	This parameter determines the values for I/O coalescing for cr_delay (msec) or cr_count outstanding commands.
lpfc_devloss_tmo	30	0	255	Seconds to hold I/O error if device disappears.

Table 13: lpfc Dynamic Parameters (Do not require a driver reload to change) (Continued)

Variable	Default	Min	Max	Comments
lpfc_enable_auth	0	0	1	This driver property specifies if the DHCHAP is enabled or not. When set to 1, DHCHAP is enabled. When set to 0, DHCHAP support is disabled. NOTE: This property requires a link reset to activate.
lpfc_fdmi_on	0	0	2	False (0) if disabled. (1) or (2) if enabled depending on type of support needed.
lpfc_link_speed	0	0=auto select 1=1 Gb/s 2=2 Gb/s 4=4 Gb/s 8=8 Gb/s		Sets link speed.
lpfc_log_verbose	0x0	0x0	0xffff	(bit mask) Extra activity logging.
lpfc_nodev_tmo (deprecated)	30	1	255	Seconds to hold I/O error if device disappears. This parameter will not work if you altered lpfc_devloss_tmo. NOTE: This is a deprecated field and lpfc_devloss_tmo should be used instead.
lpfc_pci_max_read	2048	512, 1024, 2048, 4096		Maximum DMA read byte count.
lpfc_poll	0	1= poll wiith interrupts enabled 3 = poll and disable FCP ring interrupts		Sets FCP ring polling mode control.
lpfc_poll_tmo	10	1	255	Milliseconds the driver waits between polling FCP ring interrupts.
lpfc_topology	0	0x0=loop then P2P 0x2=P2P only 0x4=loop only 0x6=P2P then loop		FC link topology (defaults to loop, if it fails attempts point-to-point mode).
lpfc_use_adisc	0	0=Off	1=On	Sends ADISC instead of PLOGI for device discovery or RSCN.

Driver for VMware ESX Configuration Parameters

All adapter-specific parameters have an `lpfcX_` prefix (where X is the driver instance number); e.g., setting `lpfc0_lun_queue_depth= 20` makes 20 the default number of maximum commands which can be sent to a single logical unit (disk) for `lpfc` instance 0.

Note: NPIV is not available on 1 Gb/s and 2 Gb/s HBAs.

Dynamic parameters do not require the driver to be unloaded and reloaded for changes to take effect.

Table 14: Driver for VMware ESX Configuration Parameters

Variable	Default	Min	Max	Dynamic	Comments
<code>lpfc_hba_queue_depth</code>	65535	0	65535	Yes	Maximum number of FCP commands that can queue to an Emulex adapter.
<code>lpfc_initiator_login</code>	0	0	1	Yes	Enables logins to other virtual initiators.
<code>lpfc_ack0</code>	0	0=Off	1=On	No	Use ACK0 for class 2.
<code>lpfc_automap</code>	1	0=Off	1=On	No	Automatically assign SCSI IDs to FCP targets detected.
<code>lpfc_check_cond_err</code>	0	0=Off	1=On	Yes	Treat certain FCP check conditions as FCP RSP errors.
<code>lpfc_cr_count</code>	1	1	255	No	This parameter determines the values for I/O coalescing for <code>cr_delay</code> (msec) or <code>cr_count</code> outstanding commands.
<code>lpfc_cr_delay</code>	0	0	63	No	This parameter determines the values for I/O coalescing for <code>cr_delay</code> (msec) or <code>cr_count</code> outstanding commands.
<code>lpfc_delay_rsp_err</code>	0	0=Off	1=On	Yes	Treat FCP RSP errors like no-device-delay.
<code>lpfc_discovery_threads</code>	1	1	64	No	Specifies the maximum number of ELS commands that can be outstanding for a discovery.
<code>lpfc_dqfull_throttle_up_inc</code>	1	0	128	Yes	Amount to increment LUN queue depth each time.
<code>lpfc_dqfull_throttle_up_time</code>	1	0	30	Yes	Time interval, in seconds, to increment LUN queue depth.
<code>lpfc_extra_io_tmo</code>	0	0	255	Yes	Extra FCP cmd timeout when connected to a fabric (in seconds).
<code>lpfc_fcp_bind_DID</code>	inactive	N/A	N/A	No	Bind specific SCSI IDs to targets based on FC Port ID.
<code>lpfc_fcp_bind_method</code>	2	1	4	No	Specifies the bind method (WWWN/WWPN/DID/ALPA map) to be used.

Table 14: Driver for VMware ESX Configuration Parameters (Continued)

Variable	Default	Min	Max	Dynamic	Comments
lpfc_fcp_bind_WWNN	inactive	N/A	N/A	No	Bind specific SCSI IDs to targets based on FC WWNN.
lpfc_fcp_bind_WWPN	inactive	N/A	N/A	No	Bind specific SCSI IDs to targets based on FC WWPN.
lpfc_fcp_class	3	2	3	Yes	FC class for FCP data transmission.
lpfc_fdmi_on	0	0	2	No	False (0) if disabled. (1) or (2) if enabled depending on type of support needed.
lpfc_inq_pqb_filter	1	0=Off	1=On	No	If true, the driver changes the peripheral quantifier bit from 1 to 3 for inquiry responses.
lpfc_iocb_watchdog_tmo	40	0	55	No	Timeout value for pending FC I/O in the driver.
lpfc_linkdown_tmo	30	0	255	Yes	(seconds) How long the driver waits before deciding that the FC link is down.
lpfc_link_speed	0	0=auto select 1=1 Gb/s 2=2 Gb/s 4=4 Gb/s		No	Sets link speed.
lpfc_log_verbose	0x0	0x0	0xffff	Yes	(bit mask) Extra activity logging.
lpfc_pci_max_read	0	0	4096	No	The maximum number of bytes transferred per pci DMA read. The default value 0 means the driver will automatically determine the correct value.
lpfc_lun_queue_depth	30	1	128	Yes	Default max commands sent to a single logical unit (disk).
lpfc_lun_skip	0	0=Off	1=On	No	Allows SCSI layers to detect all LUNs if there are LUN holes on a device.
lpfc_max_lun	256	1	256	Yes	Specifies the maximum number of LUNs per target. A value of 20 means LUNs from 0 to 19 are valid.
lpfc_max_target	256	1	256	No	This configuration parameter limits how many targets the driver can support.
lpfc_max_vpi	0xffff	0	0xffff	No	NPIV: Maximum number of vpis available per physical port.
lpfc_nodev_holdio	0	0=Off	1=On	Yes	If the device disappears, hold I/O until it comes back.

Table 14: Driver for VMware ESX Configuration Parameters (Continued)

Variable	Default	Min	Max	Dynamic	Comments
lpfc_no_device_delay	1	0	30	Yes	Delay to fail back an I/O in seconds.
lpfc_nODEV_tmo	30	0	255	Yes	Seconds to hold I/O err if device disappears.
lpfc_ns_threads	2	1	32	Yes	NPIV: Number of concurrent NameServer requests allowed to be outstanding.
lpfc_peer_vport_login	0	0	1	Yes	NPIV: Allows peer Vports to log into each other.
lpfc_scan_down	1	0=Off	1=On	Yes	Select method for scanning ALPA to assign a SCSI ID.
lpfc_scsi_req_tmo	30	0	255	Yes	Time out value (in seconds) for SCSI passthrough requests.
lpfc_tgt_queue_depth	0	0	8192	Yes	Default max commands sent to a single target.
lpfc_topology	0	0x0=loop then P2P 0x2=P2P only 0x4=loop only 0x6=P2P then loop		No	FC link topology (defaults to loop, if it fails attempts point-to-point mode).
lpfc_use_adisc	0	0=Off	1=On	Yes	Send ADISC instead of PLOGI for device discovery or RSCN.
lpfc_xmt_que_size	256	128	8192	No	Number of outstanding commands for an adapter.

Configuring Boot from SAN

You can use the HBAware utility to configure a system to boot from an attached SAN LUN. Boot from SAN allows servers on a storage network to boot their operating systems directly from a SAN storage device, typically identified by its WWPN and a LUN located on the device. By extending the server system BIOS, boot from SAN functionality is provided by the BootBIOS contained on an Emulex adapter in the server. When properly configured, the adapter then permanently directs the server to boot from a LUN on the SAN as if it were a local disk.

Boot Types

Using the Maintenance tab, you can enable, disable or configure boot from SAN for x86 BootBIOS, EFIBoot and OpenBoot (also known as FCode).

- x86 BootBIOS works with the existing BIOS on x64 and x86 systems.
- OpenBoot (FCode) works with the existing system BIOS on Solaris SPARC systems using the LPFC driver and on Linux PowerPC systems. OpenBoot is also called FCode.
- EFIBoot works with Intel Itanium 64-bit and x64-based systems and provides 64-bit system boot capability through the use of the EFI (Extensible Firmware Interface) Shell.

Emulex provides multi-platform support for boot from SAN. Universal Boot code images contain the above three types of boot code and automatically determine your system platform type and executes the proper boot code image in the adapter. These code images reside in adapter flash memory, allowing easier adapter portability and configuration between servers.

The configuration regions on the adapter store the configuration data for each of these boot types.

Note: x86 and OpenBoot share the same configuration memory space. You cannot configure an adapter for both x86 and OpenBoot *at the same time*. If you try, a message appears that the existing boot type configuration will be overwritten by the new configuration.

Note: Boot from SAN configuration does not affect current system operation. The changes only take effect upon reboot if you have configured it correctly.

Boot Device Parameters

The boot LUN for all three boot types is in the range of 0-255. EFIBoot and OpenBoot (FCode) also support an 8-byte LUN, which you can use instead of the single-byte LUN. You must select which LUN type to configure.

- For OpenBoot, you must also provide a Target ID parameter.
- The HBAnyware utility runs on a running OS, so you must boot the host to configure boot from SAN with the HBAnyware utility.
- You must work from a running host that supports the HBAnyware utility. Often, this host has booted from a direct-attached drive. With the HBAnyware utility, you can configure a direct boot host to boot from a SAN. You can modify an existing boot from SAN configuration or configure boot from SAN on an adapter for installation in another host so it can boot from SAN.
- You must know what boot-code type the adapter has; the HBAnyware utility cannot detect this. Without knowing this, you could configure a boot type but not be able to boot from it since the adapter lacks the correct boot code.
- You must know what boot type the system supports; the HBAnyware utility cannot detect this. You can configure any boot type, but if the system does not support that type, it cannot boot from SAN.
- If you manage adapters on a remote host that is running a version of the HBAnyware utility that does not support boot from SAN, the Configure Boot button does not appear.

Note: You can configure boot from SAN before boot by using the Emulex Boot BIOS setup command line interface that runs during system startup. See the Emulex Boot BIOS setup program documentation for details.

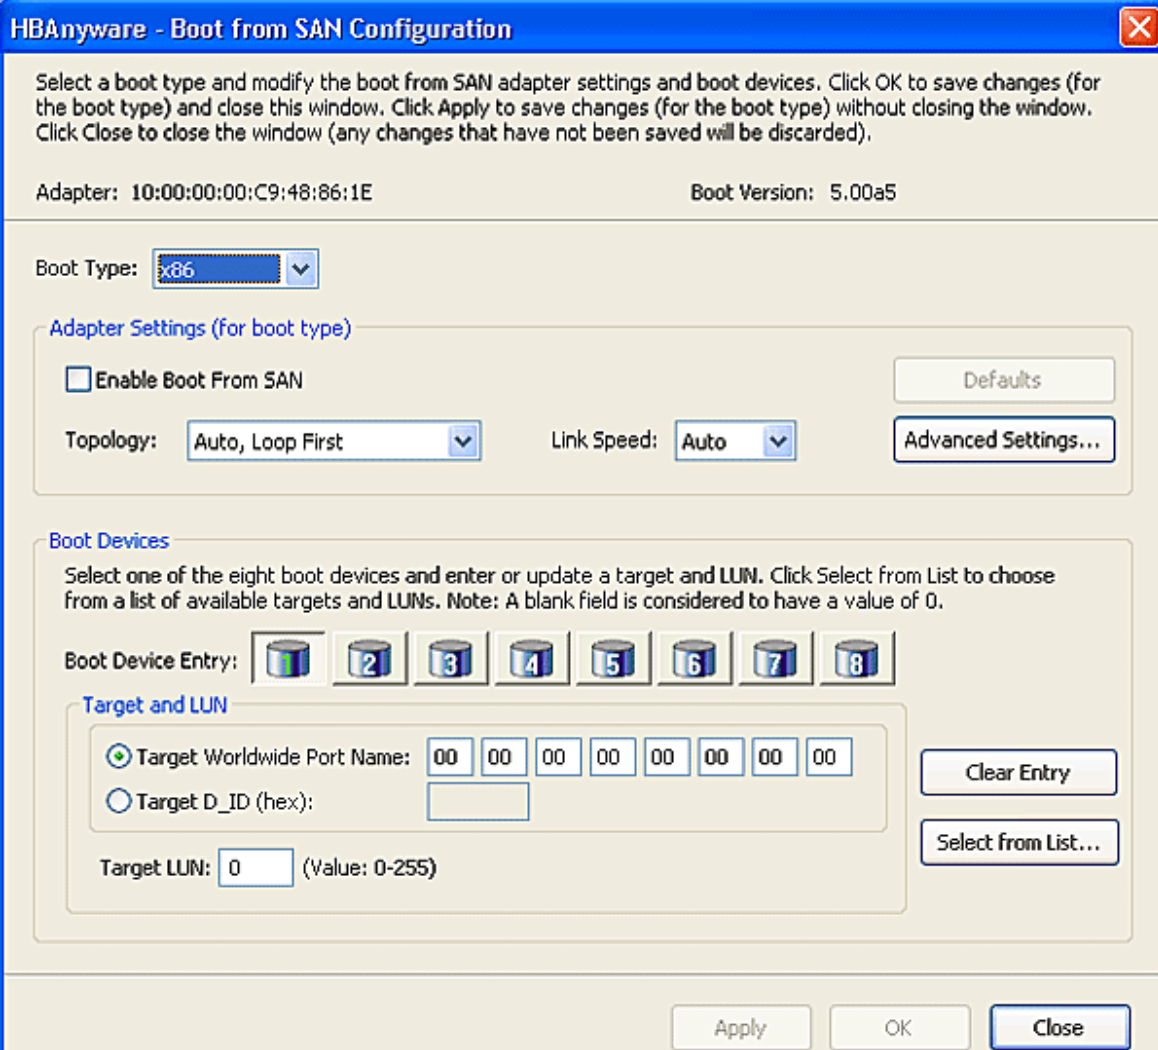
- One of the following drivers must be installed:
 - Storport Miniport for Windows
 - Emulex driver for Linux
 - Emulex LPFC driver for Solaris
 - Solaris emlxs (SFS) FCA Driver

To configure boot from SAN:

1. Select **Host View** or **Fabric View**.
2. In the discovery-tree, click the adapter port you want to enable boot from SAN.

3. Select the **Maintenance** tab, check **enable adapter boot** and click **Configure Boot**. The Boot from SAN Configuration dialog box appears.

Note: The Configure Boot button is disabled if the Enable Adapter Boot checkbox is not checked. If boot code is not present on the adapter, the Enable Adapter Boot checkbox and Configure Boot button are not displayed on the Maintenance tab.



The dialog box is titled "HBAnyware - Boot from SAN Configuration". It contains the following sections:

- Adapter:** 10:00:00:00:C9:48:86:1E
- Boot Version:** 5.00a5
- Boot Type:** A dropdown menu currently showing "x86".
- Adapter Settings (for boot type):**
 - ☐ Enable Boot From SAN
 - Topology:** A dropdown menu showing "Auto, Loop First".
 - Link Speed:** A dropdown menu showing "Auto".
 - Buttons: "Defaults" and "Advanced Settings..."
- Boot Devices:**
 - Text: "Select one of the eight boot devices and enter or update a target and LUN. Click Select from List to choose from a list of available targets and LUNs. Note: A blank field is considered to have a value of 0."
 - Boot Device Entry:** A row of eight icons representing boot devices, numbered 1 through 8.
 - Target and LUN:**
 - ☒ Target Worldwide Port Name: A field with eight "00" entries.
 - ☐ Target D_ID (hex): A single empty field.
 - Target LUN:** A field with "0" and "(Value: 0-255)".
 - Buttons: "Clear Entry" and "Select from List..."
- Buttons:** "Apply", "OK", and "Close" at the bottom right.

Figure 37: Maintenance tab

The Boot from SAN Configuration dialog box varies for each boot type. Figure 37 depicts the boot from SAN configuration for the x86 type boot.

4. Verify adapter address and boot version to make sure you configure the correct adapter and that it has the boot code version you want.
5. From the **Boot Type** menu, select x86, EFIBoot or OpenBoot.

Note: x86 and OpenBoot share the same configuration memory space. You cannot configure an adapter for both x86 and OpenBoot at the same time. When you select one of these boot types and the configuration region is configured for the other boot type, a message appears warning that making changes will overwrite the other boot-type configuration.

Note: If you have modified the settings for the current boot type and then change to a new boot type, a message appears telling you to save the current settings before changing to the new boot type.

6. Check **Enable Boot from SAN** and set the Topology and Link Speed.
 - Topology options are:
 - Auto, Loop First (default)
 - Auto, Point to Point First
 - Loop
 - Point to Point
 - Link speed options are:
 - Auto (default)
 - 1 Gb/s (if available)
 - 2 Gb/s (if available)
 - 4 Gb/s (if available)
 - 8 Gb/s (if available)
7. If you want, click Advanced Settings to configure autoscan, spinup delay and so on. See “Configuring Advanced Settings (Boot from SAN)” on page 110 for more information.
8. For x86 and EFIBoot, select one or more boot devices. For OpenBoot, select only one boot device.
9. Do one of the following on the Boot from SAN Configuration window.
 - Select **Target WorldWide Port Names**, type the numbers and click **OK**.
 - Select **Target D_ID**, type the numbers and click **OK**.
 - Select **Target LUN**, type the number and click **OK**.
 - For EFIBoot and OpenBoot, type in an 8-byte LUN (hex) and a target ID for the LUN. Also, you must enter the LUN value in “big endian” (most-significant byte, or “big end” first) order and enter all 16 characters including leading zeroes.
 - Click **Select from List**, select the target from a list of discovered LUNs (if available) and click **OK** on the Select Boot Device window. While you can manually enter the target and LUN from the Boot from SAN Configuration window, it is easier to select an existing LUN from this window. (see Figure 38) The HBAnyware utility attempts to update the boot param-

eters. If successful, a window appears with a confirmation message. Click **OK** on this confirmation window.

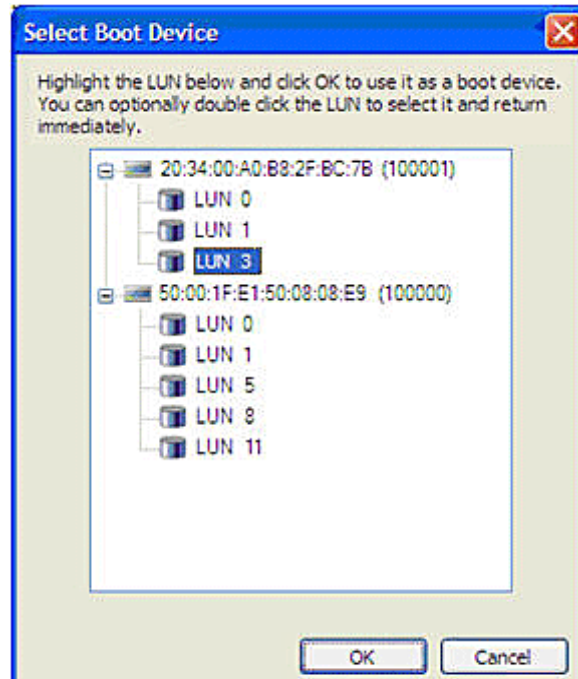


Figure 38: Select Boot Device window (for x86 or EFIBoot)

10. Click **Apply** to save your changes, but leave the dialog box open or click **OK** to apply the changes and close the dialog box.

Note: Click **Close** to close the Boot from SAN Configuration window without saving your changes. A message appears to discard your changes.

11. Reboot the system for your changes to take effect.

Configuring Advanced Settings (Boot from SAN)

The HBAware utility provides advanced settings for each boot type. From the Boot from SAN Configuration window, click **Advanced Settings**. A boot type-specific dialog box allows you to enable options such as spinup delay and autoscan. If you do not use advanced settings, the default values are used.

If you make changes you must click **OK** to save the changes and close the dialog box. You can click **Cancel** and close the dialog box without saving the changes.

Note: If you do not enter the advanced settings and the configuration for the boot type is new, default values are used. The default settings are given with descriptions of the Advanced Adapter Settings dialog boxes in the following sections.

x86 Boot Advanced Adapter Settings dialog box

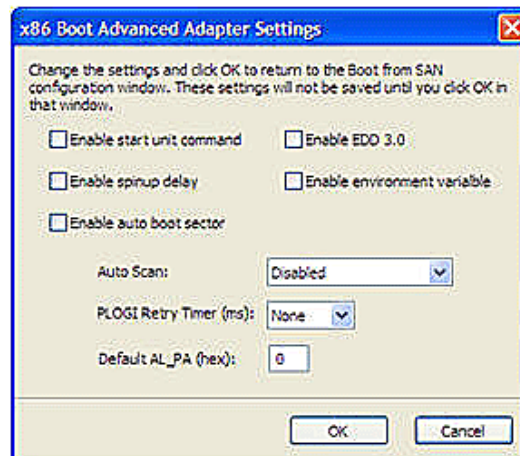


Figure 39: x86 Boot Advanced Adapter Settings dialog box

From this dialog box, configure advanced settings for the selected x86 adapter. All checkboxes are cleared (off) by default. All changes require a reboot to activate.

x86 Boot Advanced Adapter Settings Field Definitions

- Enable Start unit command - Issues the SCSI start unit command. You must know the specific LUN to issue.
- Enable EDD 3.0 - Enables the Enhanced Disk Drive (EDD) option (shows the path to the boot device). Available on Intel Itanium servers only.

Note: An x86 series system could hang during Windows 2000 Server installation if EDD 3.0 is enabled.

- Enable spinup delay - If at least one boot device has been defined, and the spinup delay is enabled, the BIOS searches for the first available boot device.
 - If a boot device is present, the BIOS boots from it immediately.
 - If a boot device is not ready, the BIOS waits for the spinup delay and, for up to three additional minutes, continues the boot scanning algorithm to find another multi-boot device.

Note: The default topology is auto topology with loop first. Change this topology setting, if necessary, before configuring boot devices.

- If no boot devices have been defined and auto scan is enabled, then the BIOS waits for five minutes before scanning for devices.
- In a private loop, the BIOS attempts to boot from the lowest target AL_PA it finds.
- In an attached fabric, the BIOS attempts to boot from the first target found in the NameServer data.
- Enable environment variable - Sets the boot controller order if the system supports the environment variable.
- Enable auto boot sector - Automatically defines the boot sector of the target disk for the migration boot process, which applies only to HP MSA1000 arrays. If there is no partition on the target, the default boot sector format is 63 sectors.
- Set Auto Scan - With auto scan enabled, the first device issues a Name Server Inquiry. The boot device is the first DID, LUN 0, or not LUN 0 device returned, depending on the option you select. Only this device is the boot device and it is the only device exported to the Multi-boot menu. Auto Scan is available only if none of the eight boot entries is configured to boot via DID or

WWPN. Emulex strongly recommends that you use the Configure Boot Devices menu to configure eight boot entries for fabric point-to-point, public loop or private loop configurations. Set to one of the following:

- Disabled (default)
- Any First Device
- First LUN 0 Device
- First non-LUN 0 Device
- Set the PLOGI Retry Timer - Sets the interval for the PLOGI (port log in) retry timer. This option is especially useful for Tachyon-based RAID arrays. Under very rare occasions, a Tachyon-based RAID array resets itself and the port goes offline temporarily in the loop. When the port comes to life, the PLOGI retry interval scans the loop to discover this device. This default setting is None (0 msec). Set to one of the following:
 - None (default)
 - 50 ms
 - 100 ms
 - 200 ms
- Type the Default AL_PA number - It has a range of 00-EF (default=0). Changes the AL_PA (Arbitrated Loop Physical Address) of the selected adapter.

EFIBoot Advanced Adapter Settings dialog box

Use the EFIBoot Advanced Adapter Settings dialog box to configure the advanced settings for the selected EFIBoot adapter.

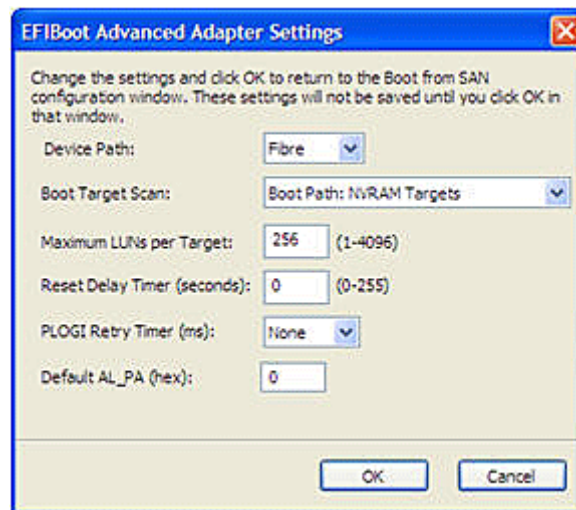


Figure 40: EFIBoot Advanced Adapter Settings dialog box

EFIBoot Advanced Adapter Settings Field Definitions

- Device Path - Makes the Fibre drive appear as a SCSI driver.
 - Fibre (default)
 - SCSI
- Boot Target Scan - This option is available only if none of the eight boot entries are configured to boot via DID or WWPN.
 - Boot Path: NVRAM Targets (default) - Discovers only LUNs that are saved to the adapter Non-Volatile Random Access Memory (NVRAM).

- Boot Path - Discovered Targets - Discovers all devices that are attached to the FC port. Discovery can take a long time on large SANs.
- None
- EFIBootFCScanLevel: NVRAM Targets and EFIBootFCScanLevel: Discovered Targets - Allows 3rd party software to toggle between Boot Path from NVRAM and Boot Path from Discovered Targets by manipulating an EFI system NVRAM variable.
- Maximum LUNs per Target - Sets the maximum number of LUNs that are polled during device discovery. The range is 1 to 4096. The default is 256.
- Reset Delay Timer in seconds - Sets a value for delay device discovery. The range is 0 to 255. The default is 0.
- PLOGI Retry Timer - Sets the interval for the PLOGI (port log in) retry timer. This option is especially useful for Tachyon-based RAID arrays. Under very rare occasions, a Tachyon-based RAID array resets itself and the port goes offline temporarily in the loop. When the port comes online again the PLOGI retry interval scans the loop to discover this device.
 - 50 ms
 - 100 ms
 - 200 ms
- Default AL_PA number – The range is 0x 00-EF. The default is 0x00. This option changes the AL_PA (Arbitrated Loop Physical Address) of the selected adapter.

OpenBoot Advanced Adapter Settings

Use this dialog box to configure the Advanced Adapter Settings for the selected OpenBoot adapter.

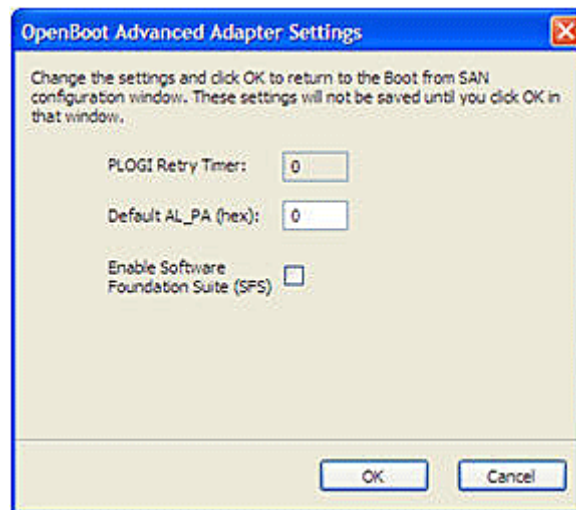


Figure 41: OpenBoot Advanced Settings dialog box

OpenBoot Advanced Adapter Field Definitions

- PLOGI Retry Timer - Sets the PLOGI Retry timer value. Range is 0 to 0xFF.
- Default AL_PA (hex) - Sets the default AL_PA. The range is 0 to 0xEF. The default is 0.
- Enable the Software Foundation Suite (SFS) - Check to enable the Software Foundation Suite (SFS) driver (the emlxs driver). The default is the lpfc driver.

Using FC-SP DHCHAP Authentication (Windows, Linux 8.2, Solaris LPFC and Solaris SFS)

Use the DHCHAP tab to view and configure FC-SP DHCHAP. You can authenticate an adapter to a switch. DHCHAP is available only for physical ports, not for virtual ports.

Note: DHCHAP is available only for physical ports, not for virtual ports.

Once DHCHAP has been activated and configured, manually initiate authentication per adapter by clicking on the Initiate Authentication button or by inducing a fabric login (FLOGI) time per the FC-SP standard to the switch. A FLOGI can also be caused by bringing the link between the switch and adapter down and then up. (Not available in read-only mode.)

Authentication must be enabled at the driver level. Authentication is disabled by default. To enable DHCHAP using the Drivers Parameters tab, enable one of the following parameters: enable-auth (in Windows), auth-mode (in Solaris LPFC), enable-auth (Solaris SFS) or lpfc-enable-auth (in Linux 8.2).

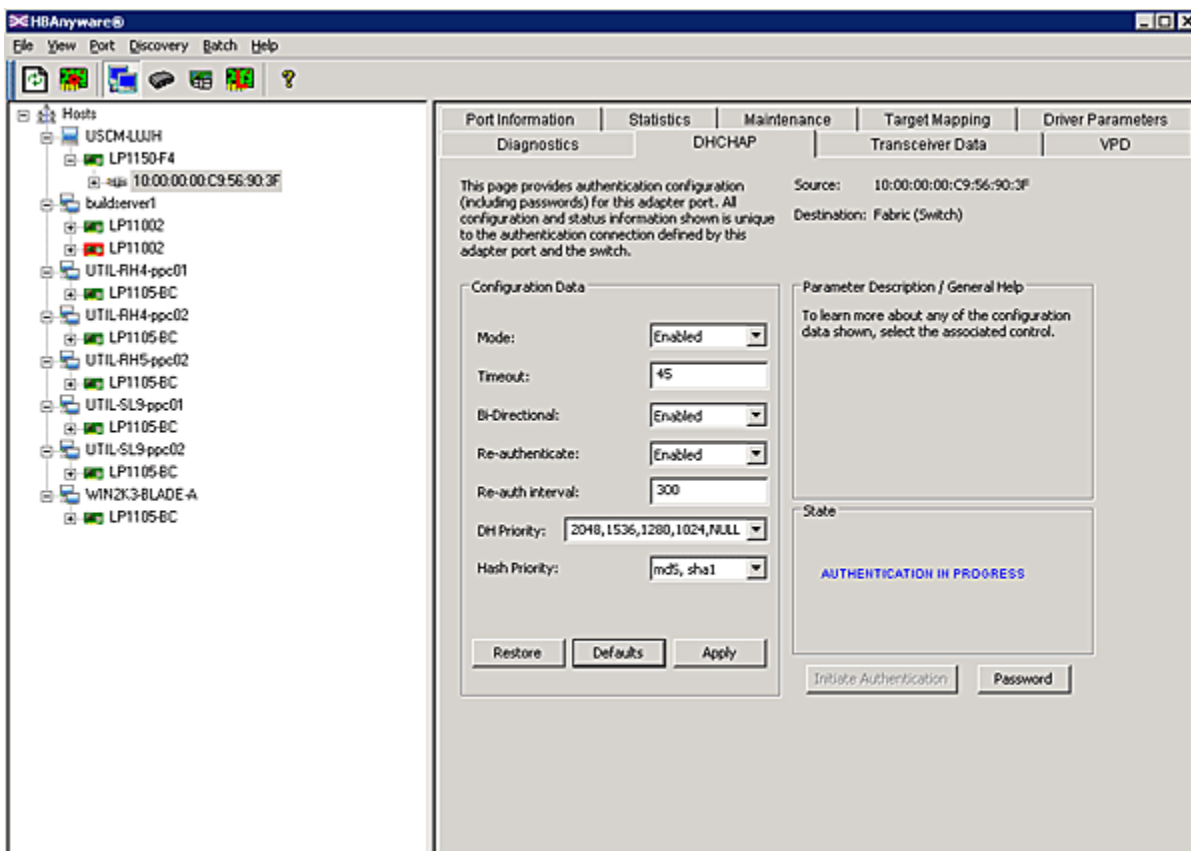


Figure 42: DHCHAP tab

DHCHAP Tab Field Definitions

- Source - The WWPN of the adapter port.
- Destination - The fabric (switch).

Configuration Data Area

- Mode - The mode of operation. There are three modes: Enabled, Passive and Disabled.

- **Enabled** - The adapter initiates authentication after issuing an FLOGI to the switch. If the connecting device does not support DHCHAP authentication, the software will still continue with the rest of the initialization sequence.
- **Passive** - The adapter does not initiate authentication, but participates in the authentication process if the connecting device initiates an authentication request.
- **Disabled** - The adapter does not initiate authentication or participate in the authentication process when initiated by a connecting device. This is the default mode.
- **Timeout** - During the DHCHAP protocol exchange, if the switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed (no authentication is performed). The time value ranges from 20 to 999 seconds.
- **Bi-Directional** - If selected, the adapter driver supports authentication initiated by either the switch or the adapter. If this checkbox is clear, the driver supports adapter initiated authentication only.
- **Re-authenticate** - If selected, the driver can periodically initiate authentication.
- **Re-authorization interval** - The value in minutes that the adapter driver uses to periodically initiate authentication. Valid interval ranges are 10 to 3600 minutes. The default is 300 minutes.
- **DH Priority** - The priority of the five supported DH Groups (Null group, and groups 1,2,3, and 4) that the adapter driver presents during the DHCHAP authentication negotiation with the switch.
- **Hash Priority** - The priority of the two supported hash algorithms (MD5 and SHA1) that the adapter driver presents during the DHCHAP authentication negotiation with the switch (default is MD5 first, then SHA1,2,3...).
- **State** - Possible states are Not Authenticated, Authentication In Progress, Authentication Success and Authentication Failed.

Changing Authentication Configuration

To view or change authentication configuration:

1. In the discovery tree, click the adapter.
2. Select the **DHCHAP** tab. (If the fields on this tab are "grayed out" (disabled) authentication has not been enabled at the driver level.)
3. Change configuration values as you want.

Note: You can only configure DHCHAP on the local host.

4. Click **Apply**. You are prompted for the current password (local password) to validate the configuration change request. The verification request only appears if a local password has been defined for this adapter.
5. Enter the password and click **OK**.

To return settings to the status before you started this procedure, click **Restore** before you click **Apply**. Once you click **Apply**, changes can not be cancelled.

To return all settings to the default configuration, click **Defaults**. Be careful as this also resets the password(s) to NULL for this configuration.

To initiate an immediate authentication, click **Initiate Authentication**. This request is sent to the driver, even if you have not made any changes to the setup.

Note: To successfully authenticate with the switch using DHCHAP, you only need to set the configuration mode to enabled and set the local password. The local password must be set to the identical value as the switch for the DHCHAP authentication to succeed.

Changing Your Password

To change your password:

1. Click **Password** on the **DHCHAP** tab. The Password dialog box is displayed.
2. Choose **Set Local Password** or **Set Remote Password**.
 - Local password is used by the adapter driver when the adapter initiates authentication to the switch (typical use).
 - Remote password is used by the adapter driver when the switch authenticates with the adapter. The latter is only possible when bi-directional has been checked on the configuration.
3. If you want to see the Password characters entered in the dialog box, check **Show Characters**.
4. Provide the current value for the password to validate the 'set new password' request (unnecessary if this is the first time the password is set for a given adapter).
5. Enter the new password.
6. Select alpha-numeric or hex format.
7. Click **OK**.

Caution: Do not forget the password once one has been assigned. Once a password is assigned to an adapter, subsequent DHCHAP configuration settings for that adapter including 'default configuration' or new passwords require you to enter the existing password to validate your request (i.e. no further changes can be made without the password).

Note: Additional help is available by clicking **Help** on the Set Password dialog box.

Viewing the Error and Event Log

For Solaris and Linux systems, a simple shell script checks the /var/log/messages files for recent Emulex driver DHCHAP events and outputs them to a default location.

To view the error and event log:

Click **Event Log History** on the **Authenticate** tab.

Updating Firmware

Updating Adapter Firmware

Using the Maintenance tab, you can update firmware on local and remote adapters. The firmware file must be downloaded from the Emulex Web site and extracted to a local drive before you can perform this procedure. (Not available in read-only mode.)

- The Emulex driver must be installed.
- The HBAnyware utility must be installed.
- The firmware zip file must be downloaded from the Emulex Web site, unzipped and extracted to a folder on a local drive.
- If the adapter is already connected to a boot device, the system must be in a state in which this type of maintenance can be performed:
 - I/O activity on the Bus has been stopped.

- Cluster software, or any other software that relies on the adapter to be available, is stopped or paused.

Note: For OEM branded HBAs, see the OEM's Web site or contact the OEM's customer service department or technical support department for the firmware files.

Note: You cannot update firmware with the HBAnyware utility on a Sun-branded adapter.

To update firmware:

1. Select **Host View** or **Fabric View**.
2. In the discovery-tree, click the adapter port whose firmware you want to update.
3. Select the **Maintenance** tab and click **Update Firmware**. The following warning screen may appear:

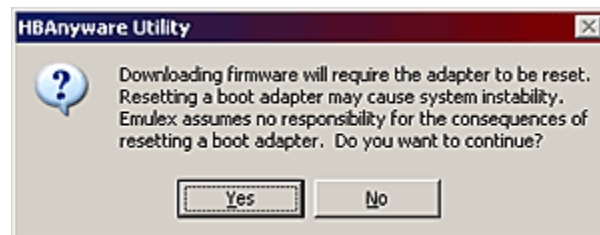


Figure 43: Reset message

4. If the warning screen appears, click **Yes**.

The Firmware Download dialog box appears (Figure 44).

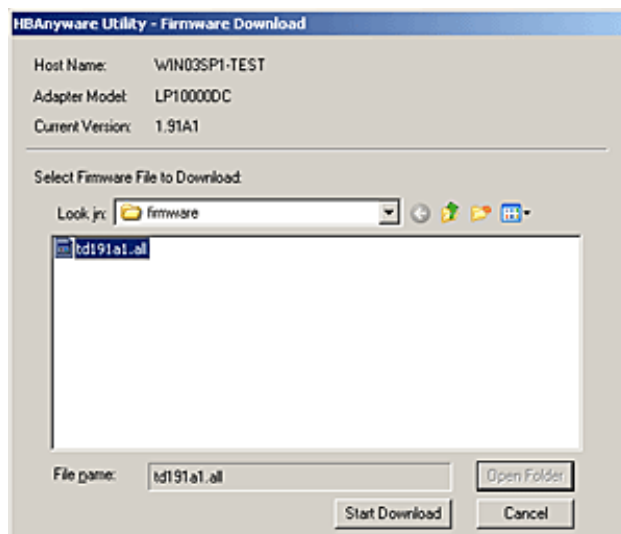


Figure 44: Firmware Download dialog box for Windows

5. Using the Firmware Download dialog box, navigate to the unzipped, extracted image file you want to download.

Note: Solaris LPFC, Solaris SFS and Linux: A Browse button is included on the Firmware Download dialog box. The Browse button and the Browse arrow allow you to navigate to a download file. Navigate to the download file and click **OK**.

6. Select the file and click **Start Download**. A status bar shows the progress of the download. The adapter in the discovery-tree is displayed in black text when the update is complete.

Note: The adapter in the discovery-tree is displayed in red text when it is offline.

7. Click **Close**. The Firmware tab displays the updated firmware information for the selected adapter.

If you are updating the firmware on a dual-channel adapter, repeat steps 1 through 6 to update the firmware on the second port or use the Updating Firmware using Batch Mode procedure.

Note: if the state of the boot code on the board has changed, this change is reflected immediately on the Port Information tab.

8. Repeat steps 2 through 6 to update boot code on additional HBAs.

Updating Adapter Firmware using Batch Mode

Use batch mode to install firmware on multiple HBAs in a single step. Batch firmware loading is restricted to a single firmware file and to all accessible HBAs for which that file is compatible. (Not available in read-only mode).

Note: Stop other HBAware utility functions while batch loading is in progress.

Before you can perform a batch update, the firmware file must be downloaded from the Emulex Web site and extracted:

- In Windows to a directory on your local drive.
- In Solaris LPFC and Solaris SFS to the Emulex Repository folder (RMRepository) in /opt/HBAware/RMRepository.
- In Linux to the Emulex Repository folder (RMRepository) in usr/sbin/hbanyware/RMRepository.

To batch load firmware:

1. From the **Batch** menu, select **Download HBA Firmware**.

Note: You do not need to select a particular tree element for this operation.

2. Windows: When the Select Firmware File dialog box appears, browse to locate and select the firmware file to download.

Solaris and Linux: Click **Browse**. A file selection dialog box appears. Select the firmware file.

3. Click **Open**.

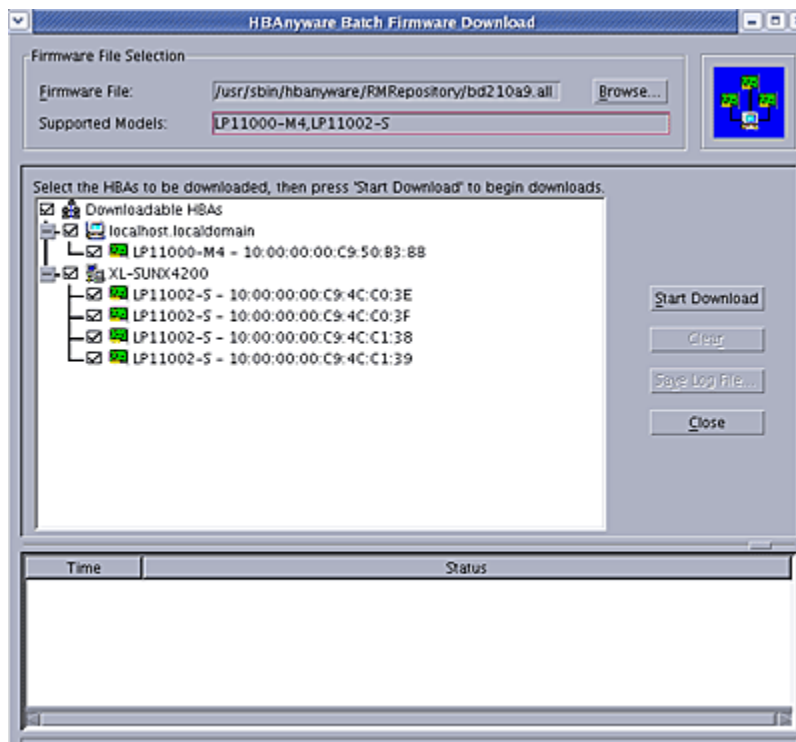


Figure 45: Selecting HBAs to Update screen

A tree-view appears showing all adapters and their corresponding hosts for which the selected firmware file is compatible. Checkboxes next to the host and adapter entries are used to select or deselect an entry. Checking an adapter selects or removes that adapter; checking a host removes or selects all eligible adapters for that host. (Figure 45).

4. Make your selections and click **Start Download**.

When downloading begins, the tree-view displays the progress. As firmware for a selected adapter is being downloaded, it appears orange in the tree-view. Once successful downloading is complete, the entry changes to green. If the download fails, the entry is changed to red.

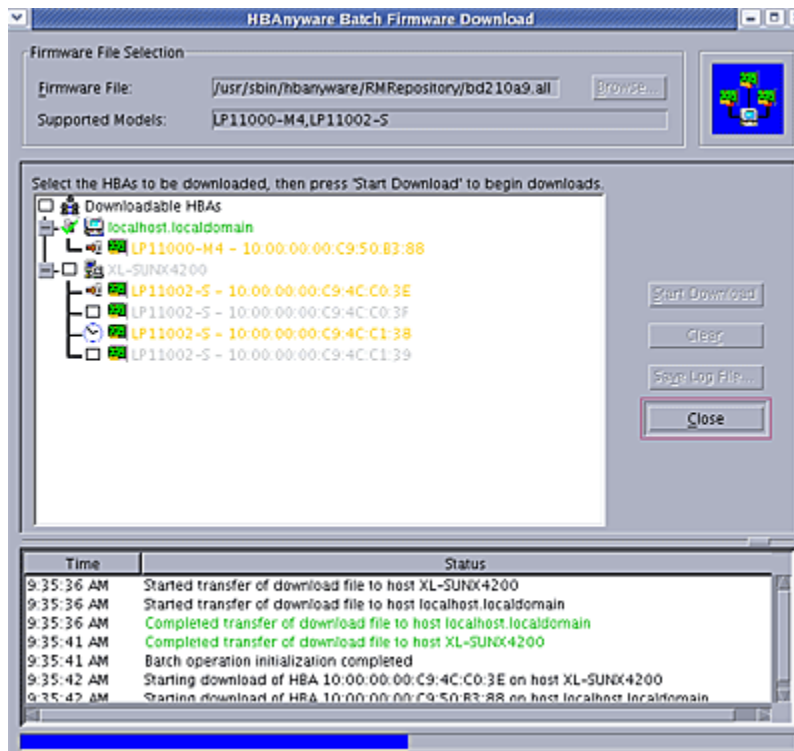


Figure 46: Download Complete screen

- When downloading is finished, you can click **Print Log** to print a hard copy of the activity log.

Note: Printing is not supported in Linux.

- Click **Close** to exit the batch procedure.

Downloading Converged Enhanced Ethernet Firmware (CEE)

To support configuration of CNAs (Converged Network Adapters) that support FCoE (Fibre Channel over Ethernet) devices, the HBAAnyware utility version 4.0 includes a CEE (Converged Enhanced Ethernet) tab. This tab is only shown when a CNA is selected in the discovery-tree. The CEE tab allows you to download firmware to the CNA port and to configure or view CEE-specific settings.

Note: CEE firmware image filenames end with a .bin extension.

To download firmware to the CNA port:

- Select **Host View** or **Fabric View**.
- In the discovery-tree, click the CNA port whose firmware you want to update.

3. Select the **CEE** tab.

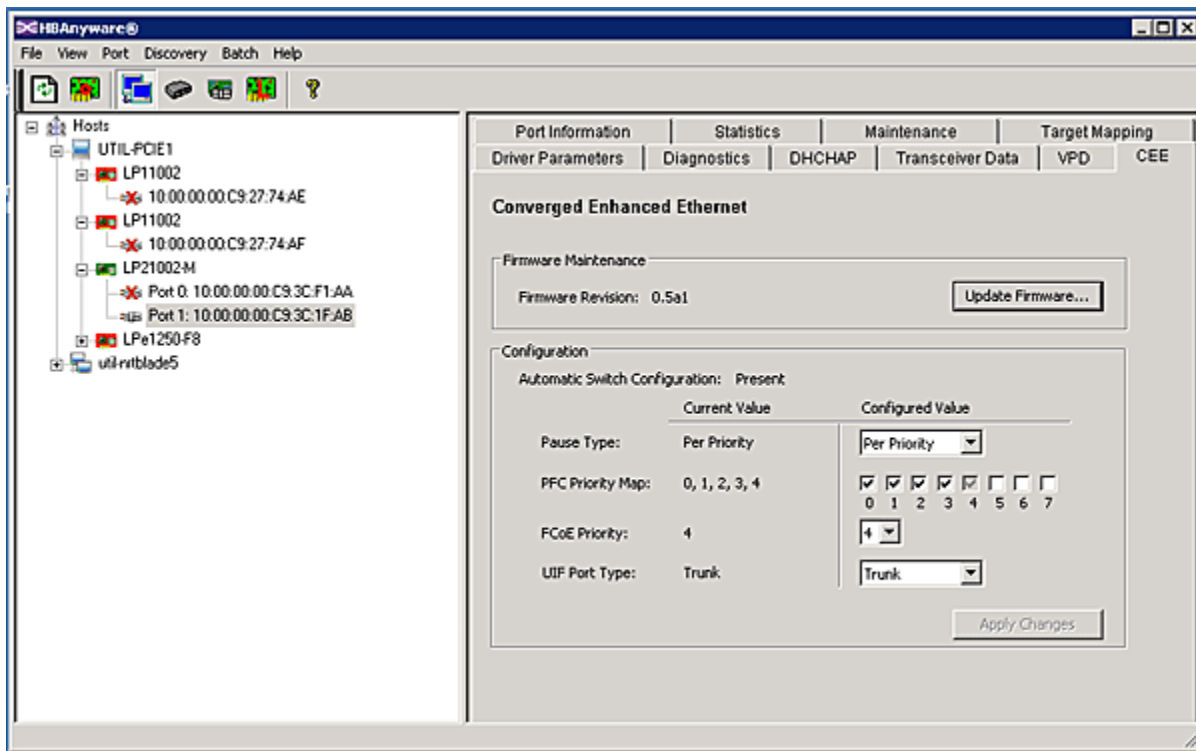


Figure 47: CEE Tab

4. Click **Update Firmware**. The CEE Firmware Download dialog box is displayed.

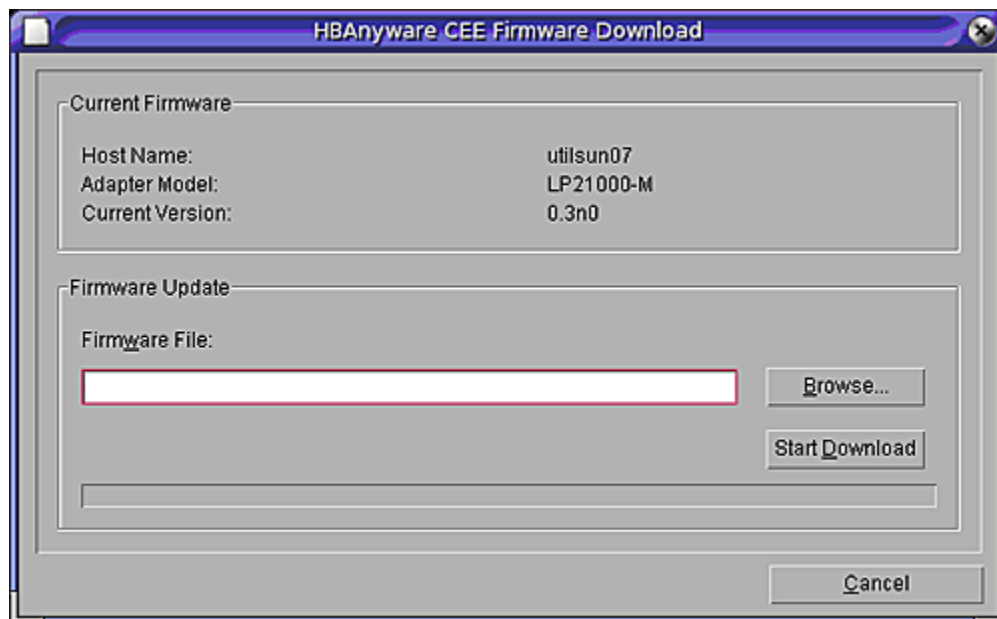


Figure 48: HBAAnyware CEE Firmware Download dialog box

5. Specify the desired firmware image. Do one of the following in the CEE Firmware Download dialog box:
 - Type the firmware file name. There are two ways to enter the file name in the Firmware File field:

- If the file is **not** located in the HBAware utility repository, type the full path and filename of the firmware image file.
- If the firmware file **is** located in the HBAware utility repository, type only the filename. The HBAware utility repository can be found in the following paths:
 - /opt/HBAware/RMRepository/ (Solaris)
 - /usr/sbin/hbanyware/RMRepository/ (Linux)
 - C:\Program Files\Emulex\Util\Emulex Repository\ (Windows)
- Click **Browse**.

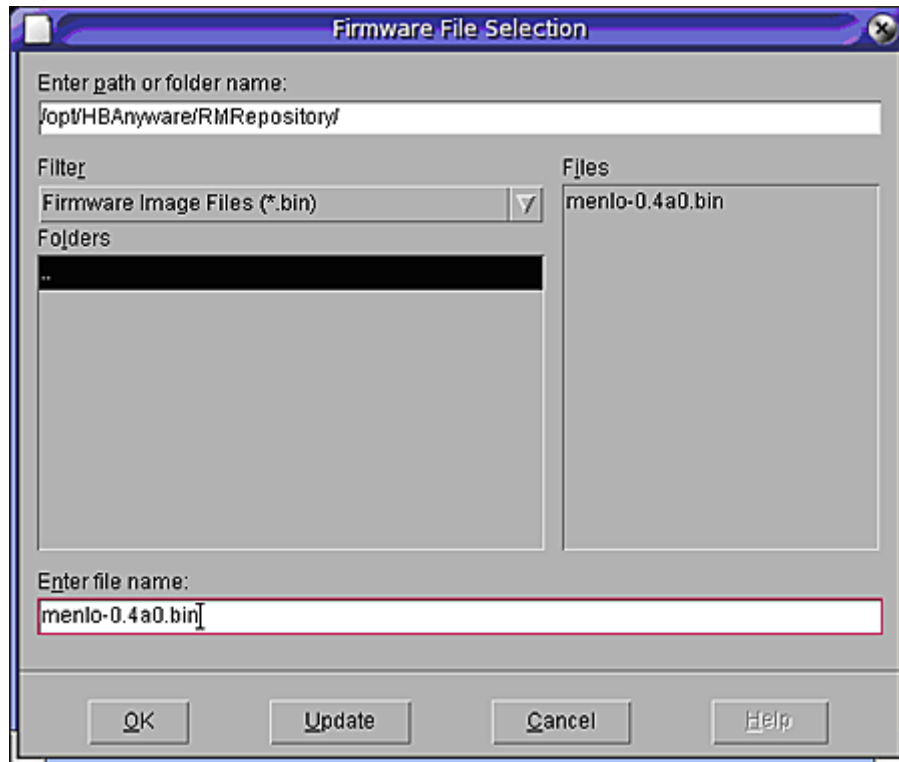


Figure 49: CEE Firmware File Selection

Use the Firmware File Selection dialog box to locate the firmware image and click **OK**. The CEE Firmware Download dialog box is displayed with the path you just browsed to.

6. Click **Start Download** on the CEE Firmware Download dialog box. A warning message similar to the following is displayed:



Figure 50: CEE Download Firmware warning

7. Click **Yes** on the Download Firmware warning and the status of the download appears on the HBAAnyware Firmware Download window, similar to the following:

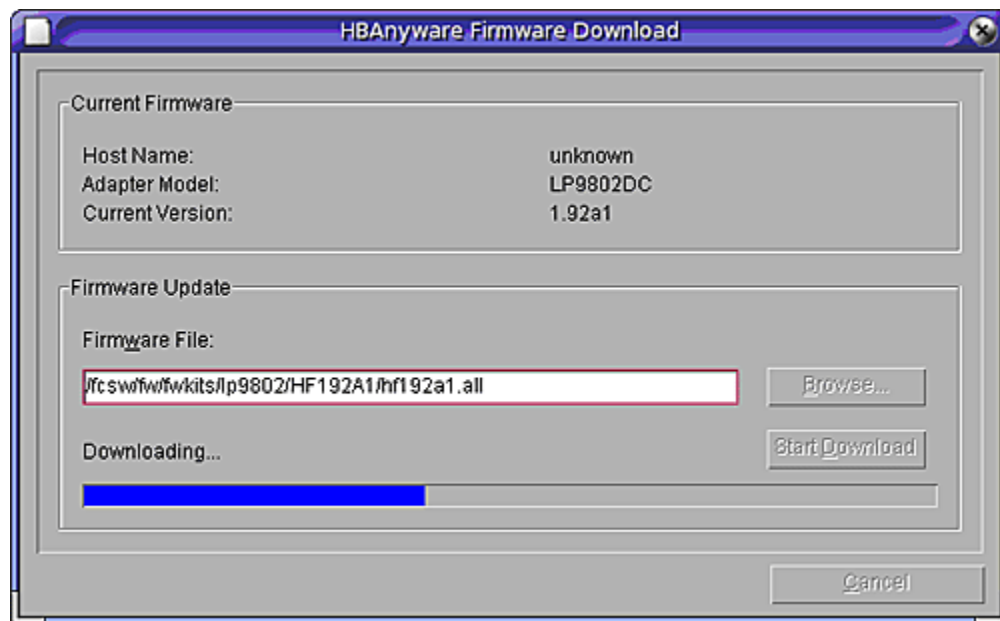


Figure 51: CEE Firmware Download status

Updating CEE Firmware using Batch Mode

Use batch mode to install CEE firmware on multiple HBAs in a single step. Batch firmware loading is restricted to a single firmware file and to all accessible HBAs for which that file is compatible. (Not available in read-only mode).

Note: Stop other HBAware utility functions while batch loading is in progress.

Before you can perform a batch update, the firmware file must be downloaded from the Emulex Web site and extracted:

- In Windows to a directory on your local drive.
- In Linux to the Emulex Repository folder (RMRepository) in /usr/sbin/hbanyware/RMRepository.
- In Solaris LPFC and Solaris SFS to the Emulex Repository folder (RMRepository) in /opt/HBAware/RMRepository.

To batch load CEE firmware:

1. From the **Batch** menu, select **Download CEE Firmware**. The following warning is displayed:

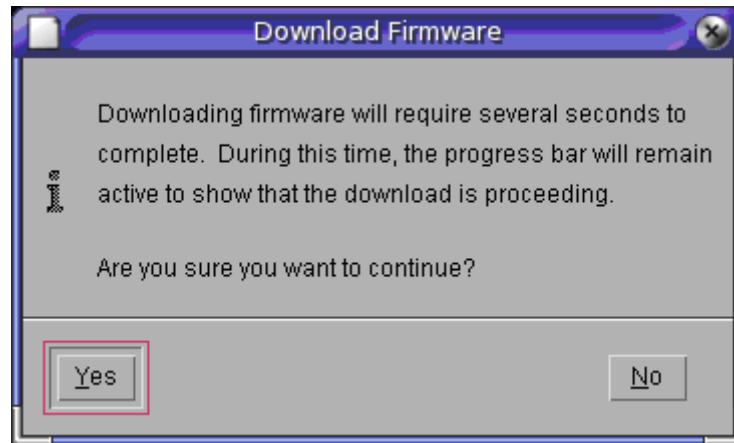


Figure 52: Download CEE Firmware warning

Note: You do not need to select a particular tree element for this operation.

2. Click **Yes**. The Select Firmware File dialog is displayed.

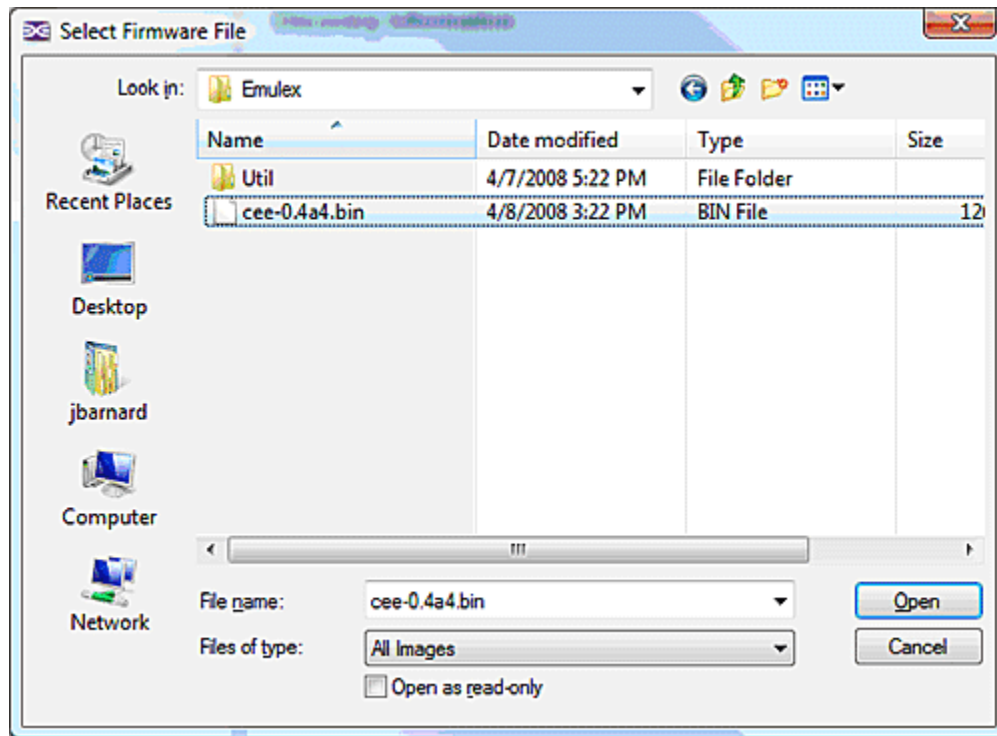


Figure 53: Cee Firmware File Selection dialog

3. Windows: When the Select Firmware File dialog box appears, browse to locate and select the firmware file to download.

Solaris and Linux: Click **Browse**. A file selection dialog box appears. Select the firmware file.

- Click **Open**.

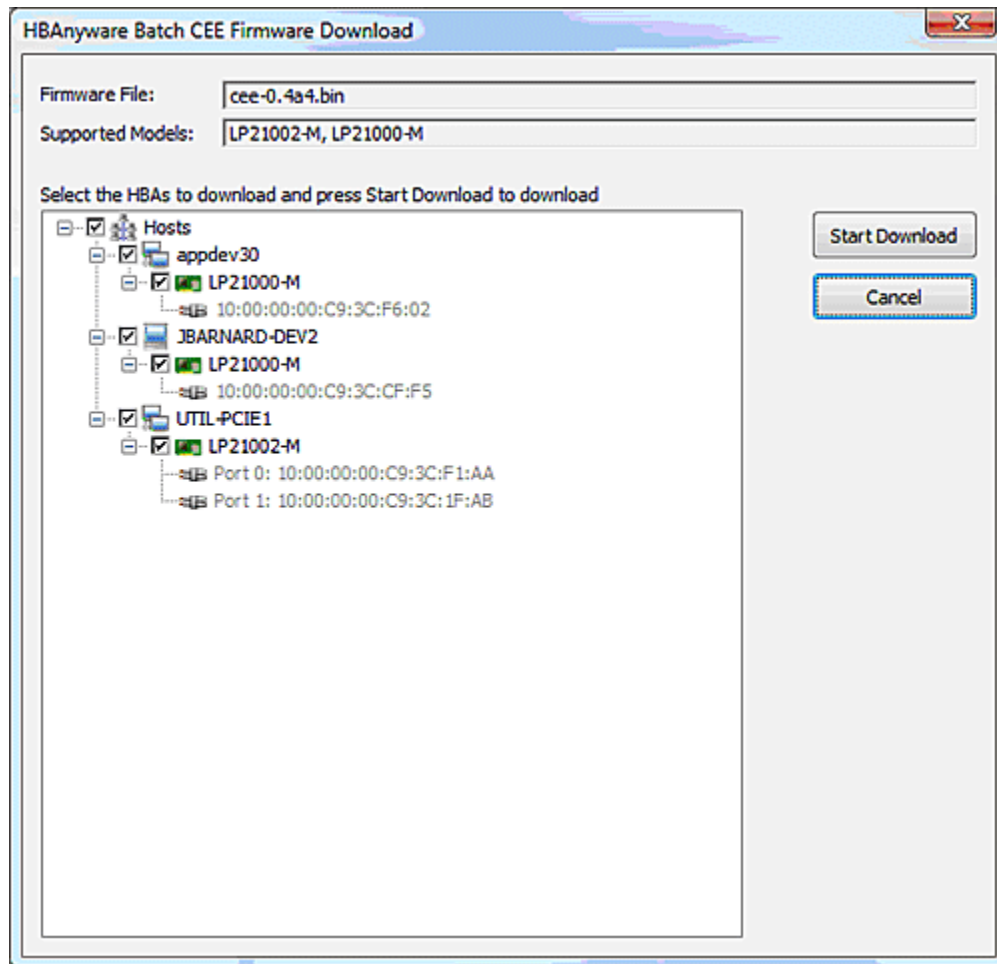


Figure 54: Selecting HBAs to Update screen

A tree-view appears showing all adapters and their corresponding hosts for which the selected firmware file is compatible. Checkboxes next to the host and adapter entries are used to select or deselect an entry. Checking an adapter selects or removes that adapter; checking a host removes or selects all eligible adapters for that host. (Figure 45).

- Make your selections and click **Start Download**.

When downloading begins, the tree-view displays the progress. As firmware for a selected adapter is being downloaded, it appears orange in the tree-view. Once successful downloading is complete, the entry changes to green. If the download fails, the entry is changed to red.

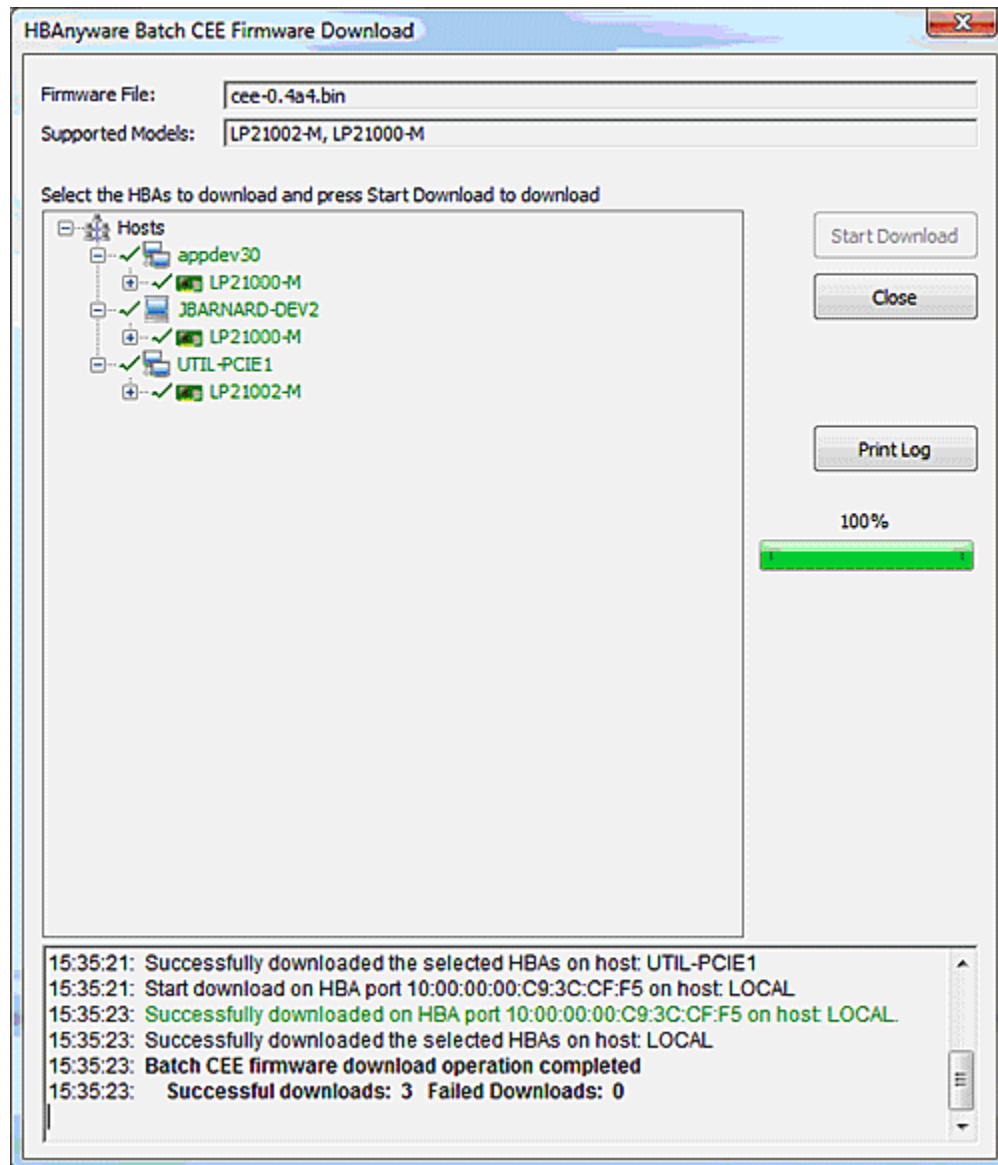


Figure 55: Download Complete screen

6. When downloading is finished, you can click **Print Log** to print a hard copy of the activity log.

Note: Printing is not supported in Linux.

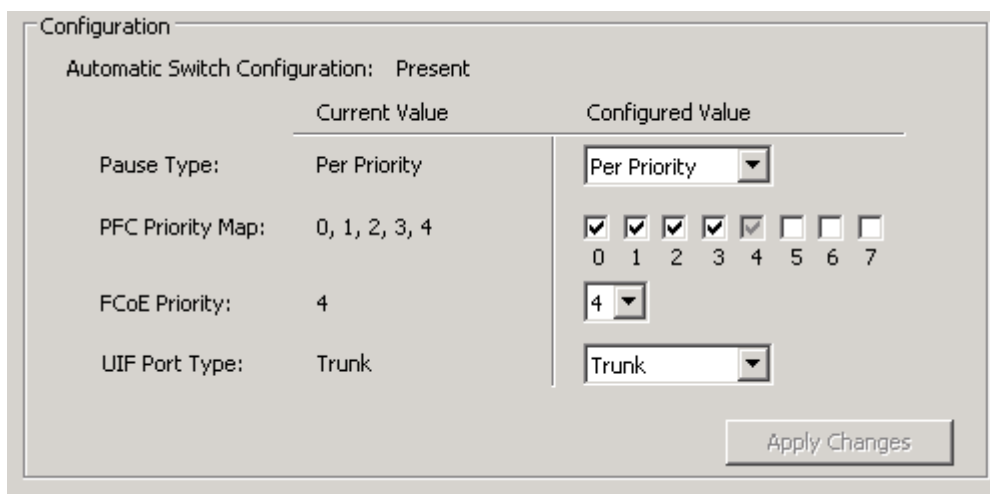
7. Click **Close** to exit the batch procedure.

Configuring CEE-Specific Parameters

The CEE (Converged Enhanced Ethernet) tab appears only if you select a CEE adapter (such as an LP21000) from the discovery-tree and if the CEE tab is enabled (See “Customizing Tab Views” on page 23). The CEE tab allows you to view and configure CEE-specific parameters for the selected port.

- If the port is connected to a fabric switch with Automatic Switch Configuration present, the Current Value column will display the settings currently being used by the switch. Any changes in the Configured Value column will be saved to non-volatile memory, but will not go into effect until the port is connected to a switch without Automatic Switch Configuration present.
- If the port is connected to a switch without Automatic Switch Configuration present, the Current Value column will show the values currently in use by the CEE adapter port. When changes are applied to the Configured Value column, the CEE adapter will attempt to use these settings, but the switch may not allow it (depending on it's configuration). Changes that are successfully loaded will appear in the Current Value column.

Note: The CEE-specific parameters cannot be set when the link is down.



The screenshot shows a 'Configuration' window with a title bar. Inside, there's a section 'Automatic Switch Configuration: Present'. Below this is a table with two columns: 'Current Value' and 'Configured Value'. The rows are: 'Pause Type' (Current: Per Priority, Configured: Per Priority dropdown), 'PFC Priority Map' (Current: 0, 1, 2, 3, 4, Configured: checkboxes for 0-7, with 0-4 checked), 'FCoE Priority' (Current: 4, Configured: 4 dropdown), and 'UIF Port Type' (Current: Trunk, Configured: Trunk dropdown). An 'Apply Changes' button is at the bottom right.

Figure 56: CEE tab, Configuration Data area

Configuration Data Area Field Definitions

- Automatic Switch Configuration - A non-configurable field displays whether the Automatic Switch Configuration feature is present on the attached switch. Possible values are Present and Not Present.
- Pause Type - Select the Ethernet flow control type. Select between standard PAUSE flow control and Per Priority based PAUSE flow control. Per Priority based flow control means the Ethernet network is seen as 8 virtual lanes (a.k.a. “Priorities”) of traffic rather than one. Possible drop down values are Standard and Per Priority.
- PFC Priority Map - A series of eight checkboxes that can only be selected if the Pause Type is set to “Per Priority”. Selected values correspond to the flow control priorities being used by the board. The value of the FCoE Priority must always be included among the PFC Priority Map values. Select a number of values from 1 to 8. Possible values are 0 to 7.
- FCoE Priority - The available values for the FCoE Priority parameter. Possible drop down values are 0 to 7.
- UIF Port Type - Select between Access and Trunk port types.

Configuration Area Buttons

- **Apply Changes** - Applies any changes made under the Configured Value column. If Automatic Switch Configuration is present on the attached fabric switch, these changes will be saved in non-volatile memory, but not loaded. If Automatic Switch Configuration is not present, changes made in the Configured Value column may or may not take effect, depending on the switch's configuration. You will be notified of any failures to save the configured values to the CEE adapter's non-volatile memory.

Exporting SAN Information

The HBAware utility enables you to create reports about discovered SAN elements. Reports are generated in .xml and .csv format and include all the SAN information that is displayed through the various HBAware utility tabs.

Note: Creating a SAN report can take several minutes for a large SAN.

To create a SAN report:

1. From the **File** menu, select **Export SAN Info**.
2. Browse to a folder and enter a filename with .xml or .csv extension.
3. Click **Save** to start the export process.

During the export process, progress is displayed in the lower right hand side of the progress bar. On Windows, you cannot change views, reset, or download firmware during the export process.

Mapping and Masking

Automapping SCSI Devices (Windows)

The driver defaults to automatically mapping SCSI devices. The procedures in this section apply if the default has been changed.

To automap SCSI devices:

1. Display driver parameters for the host or adapter - select the **Driver Parameters** tab or the host **Driver Parameters** tab.
2. Select the **AutoMap** parameter. Several fields about the parameter appear on the right side of the tab.
3. Select **Enabled**.
4. To apply your changes, click **Apply**.
5. Reboot the system for this change to take effect.

Mapping and Masking Defaults (Windows)

Table 15: Mapping and Masking Window Defaults

Field (Function)	Default	Description	Window
Globally Automap All Targets	Enabled	Emulex driver detects all FC devices attached to the Emulex adapters.	Global Automap
Globally Automap All LUNs	Enabled	Assigns an operating system LUN ID to a FC LUN ID for all LUNs behind all targets in the system area network.	Global Automap
Globally Unmask All LUNs	Enabled	Allows the operating system to see all LUNs behind all targets.	Global Automap
Automap All LUNs (Target Level)	Disabled	With Globally Automap All LUNs disabled, this parameter assigns an operating system LUN ID to a FC LUN ID for all LUNs behind the selected target.	LUN Mapping
LUN Unmasking (Target Level)	Disabled	Allows the operating system to see all LUNs behind the selected target. With this parameter disabled, each individual LUN can be masked or unmasked.	LUN Mapping

Masking and Unmasking LUNs (Windows, Solaris LPFC and Solaris SFS)

LUN masking refers to whether or not a LUN is visible to the operating system. A LUN that has been masked is not available and is not visible to the OS. You can use the HBAnyware utility to mask or unmask LUNs at the host level.

Note: The LUN Masking tab is not shown in Virtual Port view because LUN masking is not available for virtual ports.

Note: In Solaris systems, the Emulex LPFC drivers support both a target level and adapter level LUN unmasking override feature. If either of these driver-specific overrides are enabled, the HBAnyware utility does not permit you to configure LUN masking. In this case you must change the LUN masking level to the correct level from the LUN masking tab before you can mask or unmask LUNs (see Figure 57).

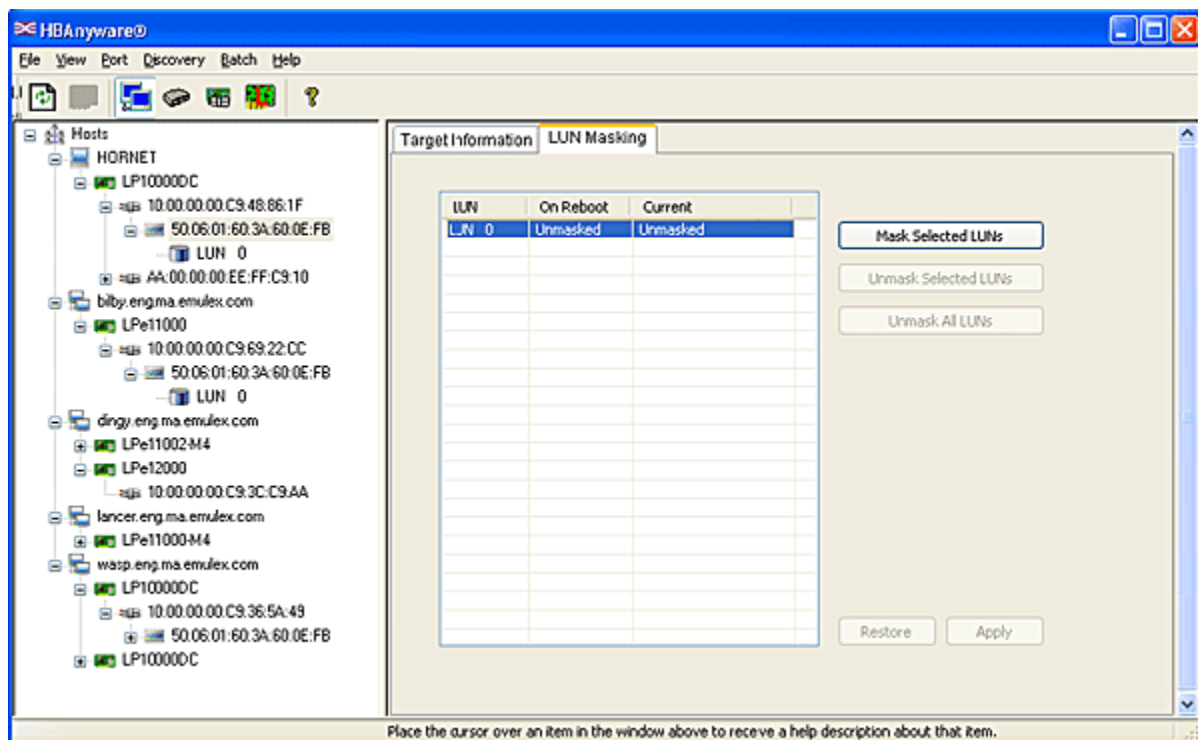


Figure 57: LUN Masking tab

LUN Masking Conventions and Guidelines

LUN icons in the discovery-tree reflect the live mask state currently in use by the driver. Green LUN icons indicate unmasked LUNs. Gray LUN icons indicate masked LUNs. Red text indicates that a LUN mask has been changed, but not applied (saved).

LUN Masking Column Definitions

- **LUN** – The FC LUN number.
- **On Reboot** – The 'On Reboot' column shows the mask configuration currently saved to the configuration file on disk (Solaris LPFC and Solaris SFS) or to the Registry (Windows). Normally, for a specific LUN, the states reported in the 'On Reboot' and 'Current' column are identical. However, there can be times where these do not match. For example, the hbacmd tool can be used to change only the 'Current' mask state for a LUN and not touch the 'On Reboot' mask state contained in the configuration file.
- **Current** – The 'Current' column displays the live mask state currently in use by the driver. When you first see the LUN Masking tab, the mask states displayed in the 'Current' column are identical to the mask states for the corresponding LUNs in the discovery-tree.

To change the mask status of a LUN:

1. Select **Host View**.
2. From the discovery-tree, select the SCSI target whose LUN masking state you want to change. A set of LUNs appears below the selected SCSI target.
3. Select the **LUN Masking** tab. This tab contains a list of the same set of LUNs that appear below the SCSI target in the discovery-tree.

4. In the LUN list of the LUN Masking tab, select one or more LUNs. The Mask Selected LUNs, Unmask Selected LUNs, Unmask All LUNs, Restore and Apply buttons become active as appropriate. For example, if the LUN is currently unmasked, only the Mask Selected LUNs button is active.
5. Change the mask status: click **Mask Selected LUN(s)**, **Unmask Selected LUN(s)** or **Unmask All LUNs** as appropriate. Mask status changes appear in red text.

Note: To return all mask settings to their status before you started this procedure, click Restore before you click Apply. Once you click Apply, changes cannot be cancelled by clicking Restore. To unmask all LUNs, click Unmask All LUNs. This button is always active. Be sure to also click Apply to commit the changes.

6. Click **Apply** to commit the changes. An informational message is displayed that confirms the mask status has changed and the red text changes to black.

Using Automapping and Persistent Binding (Windows, Solaris LPFC and Solaris SFS)

Set up persistent binding on remote and local adapters. Global automapping assigns a binding type, target ID, SCSI Bus and SCSI ID to the device. The binding type, SCSI Bus and SCSI ID can change when the system is rebooted. With persistent binding applied to one of these targets, the WWPN, SCSI Bus and SCSI ID remain the same when the system is rebooted. (Not available in read-only mode.)

The driver refers to the binding information at bootup. When you create a persistent binding, the HBAnyware utility tries to make that binding dynamic. However, the binding must meet all of the following criteria to be dynamic:

- The SCSI ID (target/bus combination) specified in the binding request must not be mapped to another target. For example, the SCSI ID must not already appear in the 'Current Mappings' table under 'SCSI ID'. If the SCSI ID is already in use, then the binding cannot be made dynamic, and a reboot is required.
- The target (WWPN, WWNN or DID) specified in the binding request must not be mapped to a SCSI ID. If the desired target is already mapped, then a reboot is required.
- The bind type (WWPN, WWNN or DID) specified in the binding request must match the currently active bind type shown in the Current Settings area of the Target Mapping tab. If they do not match, then the binding cannot be made active.

Changing Automapping Settings

To change automapping settings:

1. Select **Host View** or **Fabric View**.
2. In the discovery-tree, select the adapter port you want to set up with persistent binding.
3. Select the **Target Mapping** tab. All targets are displayed.

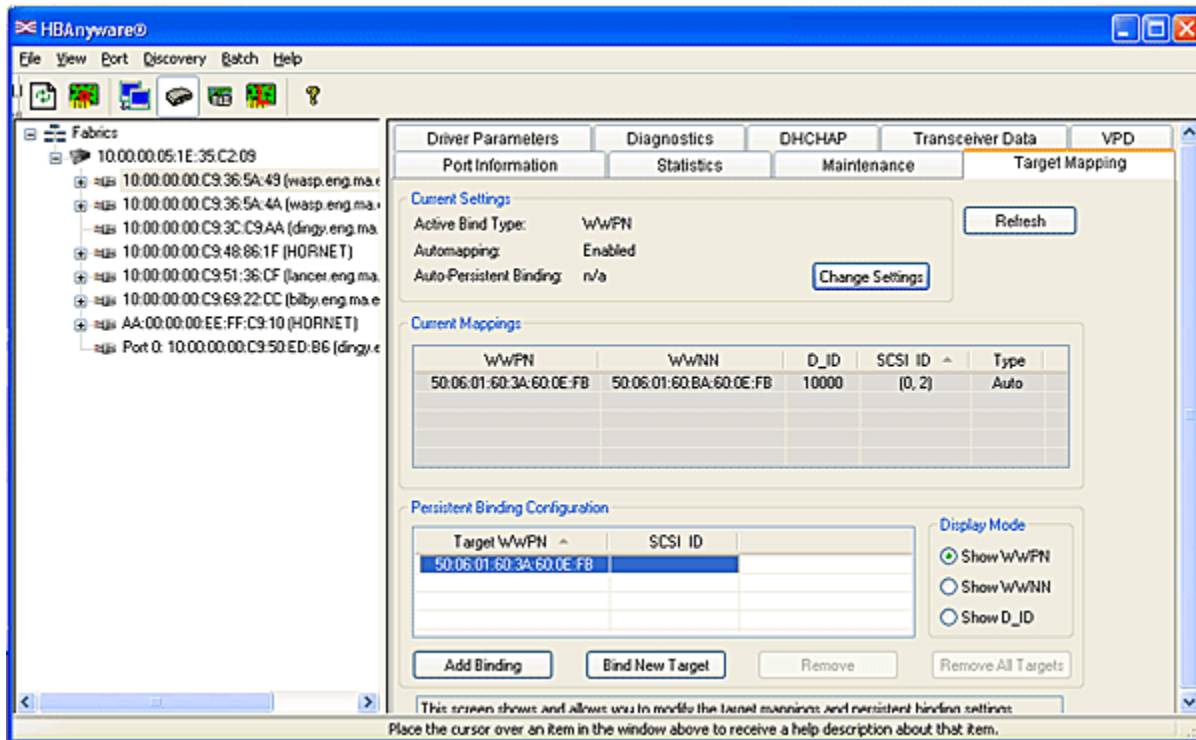


Figure 58: Target Mapping tab

4. Target mappings are displayed by WWPN, WWNN, or D_ID. "PB", indicates mapping from persistent binding, while "Auto", indicates an automapped target. In the Display Mode section, choose the display mode you want to use.
5. If you want click **Change Settings**. The Mapped Target Settings dialog box appears. You can enable or disable auto-mapping and change the active bind type. Click **OK**.
6. Reboot the system for changes to take effect.

Adding a Persistent Binding

To add a persistent binding:

1. Select **Host View** or **Fabric View**.
2. In the discovery-tree, select the adapter port you want to set up with persistent binding.
3. Select the **Target Mapping** tab. All targets are displayed. In the Targets Table, click the target that you want to bind.

- Click **Add Binding**. The Add Persistent Binding dialog box is displayed.

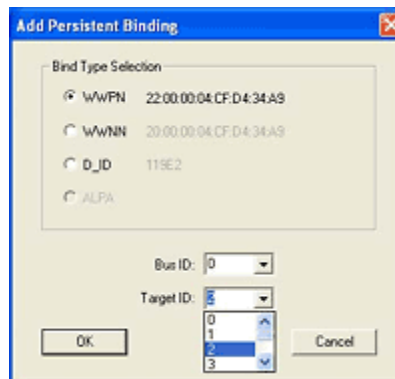


Figure 59: Add Persistent Binding dialog box

- Select the bind type that you want to use (WWPN, WWNN or D_ID).
- Select the Bus ID and target ID that you want to bind, and click **OK**.

Note: Automapped targets have entries only in the second column of the Targets Table. Persistently bound targets have entries in the second and third columns. In this case, the third column contains the SCSI Bus and target numbers you specified in the Add Persistent Binding dialog box. This binding takes effect only after the local machine is rebooted.

Binding a Target that Does Not Appear in the Persistent Binding Table

To bind a target that does not appear in the Persistent Binding table on the Target Mapping tab:

Note: It is possible to specify a SCSI Bus and target that have already been used on behalf of a different FC target. Attempting to bind a target already in the Persistent Binding table on the Target Mapping tab results in an error message, "Target already in target list. Use the Add Binding button."

- Select **Host View** or **Fabric View**.
- In the discovery-tree, select the adapter port you want to set up with persistent binding.
- Select the **Target Mapping** tab. All targets are displayed.
- Click **Bind New**. The Bind New Target dialog box is displayed.

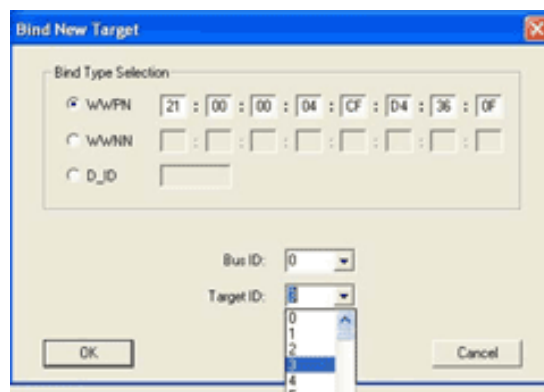


Figure 60: Bind New Target dialog box

- Click the type of binding you want to use, and type the WWPN, WWNN or D_ID you want to bind to the target.

6. Select the Bus ID and Target ID that you want to bind, and click **OK**.

Note: A target does not appear on the target list if automapping is disabled and the target is not already persistently bound.

Adding New Targets Using sd.conf for Solaris 8, 9 and 10

You can perform on-the-fly configuration changes, without rebooting, using the HBAnyware utility. For Solaris 8, you must first add the new targets to the sd.conf file.

To add new targets using sd.conf (Solaris 8):

1. Edit the Solaris SCSI configuration file (sd.conf):

```
#vi /kernel/drv/sd.conf  
  
.  
.  
.  
name="sd" parent="lpfc" target=17 lun=1;  
name="sd" parent="lpfc" target=18 lun=10;  
name="sd" parent="lpfc" target=19 lun=15;  
.  
.  
.
```

2. Save the file and exit the text editor.

Diagnostics

Note: Diagnostics are not available on VMware ESX Server.

Note: Quick Test, POST Test, and the Advanced Diagnostics Test buttons are disabled for any remote adapter that is managed in-band.

Use the Diagnostics tab to do the following:

- View flash load list, PCI registers and wakeup parameter information
- Run these tests on Emulex adapter's installed in the system:
 - PCI Loopback (see page 141) (Not available in read-only mode.)
 - Internal Loopback (see page 141) (Not available in read-only mode.)
 - External Loopback (see page 141) (Not available in read-only mode.)
 - Power-On Self Test (POST) (see page 138) (Not available in read-only mode.)
 - Echo (End-to-End) (see page 142) (Not available in read-only mode.)
 - Quick Test (see page 137) (Not available in read-only mode.)
- Perform a diagnostic dump (see page 138) (Not available in read-only mode.)
- Control adapter beaconing (see page 138) (Not available in read-only mode.)

All functions are supported locally and remotely.

Viewing Flash Contents, PCI Registers and Wakeup Information

The Diagnostic tab shows PCI register dump information and flash memory contents. The information is read-only and is depicted below:

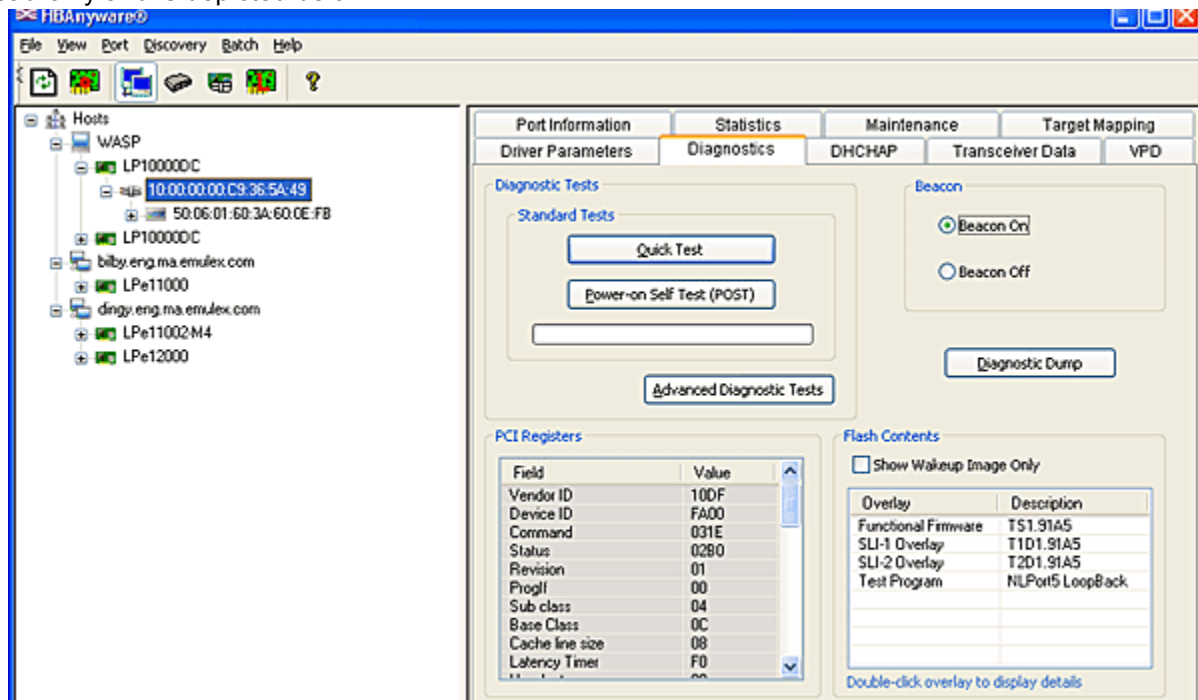


Figure 61: PCI Registers and Flash Contents of the Diagnostics tab

Viewing Flash Contents

If you check the **Show Wakeup Images Only** checkbox, the flash overlays that are not loaded when the system is booted no longer display. This checkbox defaults to unchecked.

Viewing Overlay Details

If you double-click on a flash overlay, another window appears with details about that overlay.

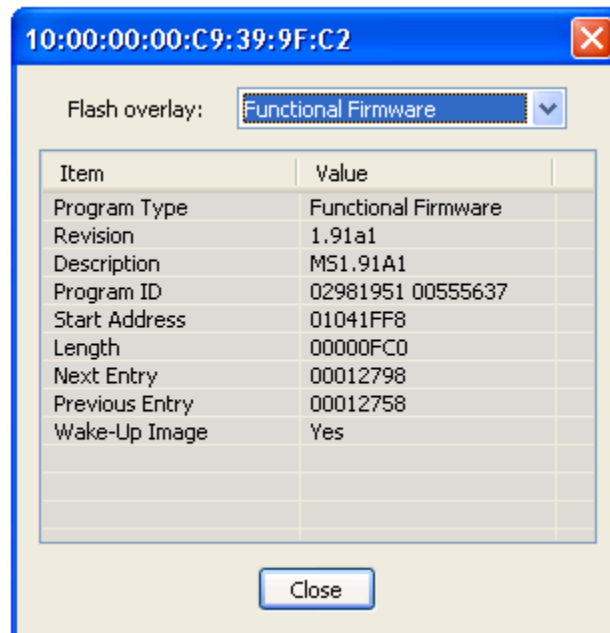


Figure 62: Overlay Detail window

To see the details of a different flash overlay image, you can either close the details window and double-click on another overlay name, or choose a different overlay name from the Flash overlay menu.

Note: The internal loopback tests are not run on the LP21000 and LP21002 adapters.

Viewing the PCI Registers

The PCI Registers appear directly on the Diagnostics tab.

Running a Quick Test

The Diagnostics tab enables you to run a “quick” diagnostics test on a selected adapter. The Quick Test consists of 50 PCI Loopback test cycles and 50 Internal Loopback test cycles. (Not available in read-only mode.)

Note: Internal and external loopback tests are not available for LP21000 and LP21002 adapters.

To use quick test:

1. From the discovery-tree, select the adapter port on which you want to run the Quick Test.
2. Select the **Diagnostics** tab and click **Quick Test**. A warning message appears.

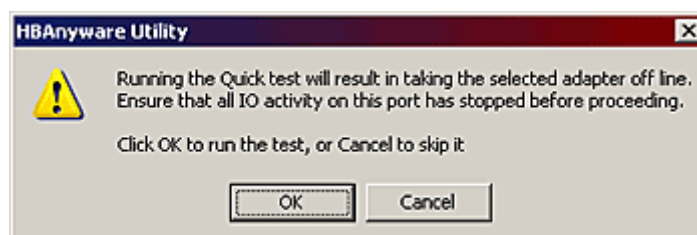


Figure 63: Quick Test Warning window

3. Click **OK** to run the test. The Quick Diagnostic Test window appears displaying the PCI Loopback and Internal Loopback test results.

Running a Power On Self Test (POST)

The POST is a firmware test normally performed on an adapter after a reset or restart. The POST does not require any configuration to run. (Not available in read-only mode.)

To run the POST:

1. From the discovery-tree, select the adapter port on which you want to run the POST.
2. Select the **Diagnostics** tab and click **Power-on Self Test (POST)**. A warning dialog box appears.
3. Click **OK**. A POST window appears displaying POST information.

Using Beaconing

The beaoning feature enables you to force a specific adapter's LEDs to blink in a particular sequence. The blinking pattern acts as a beacon, making it easier to locate a specific adapter among racks of other adapters. (Not available in read-only mode.)

When you enable beaoning, the two LEDs blink rapidly in unison for 24 seconds, after which the LEDs report the adapter health status for 8 seconds. When the 8 seconds are up, the adapter returns to beaoning mode. This cycle repeats indefinitely until you disable this feature or you reset the adapter.

Note: The beaoning buttons are disabled if the selected adapter does not support beaoning.

To enable or disable beaoning:

1. From the discovery-tree, select the adapter port whose LEDs you want to set.
2. Select the **Diagnostics** tab and click **Beacon On** or **Beacon Off**.

Creating Diagnostic Dumps

The diagnostic dump feature enables you to create a “dump” file for a selected adapter. Dump files contain various information such as firmware version, driver version and so on, that is particularly useful when troubleshooting an adapter. (Not available in read-only mode.)

Caution: Disruption of service can occur if a diagnostic dump is run during I/O activity.

To start a diagnostic dump:

1. From the discovery-tree, select an adapter port whose diagnostic information you want to dump.

2. Select the **Diagnostics** tab and click **Diagnostic Dump**. The Diagnostic Dump dialog box appears. You can specify how many files you want to save using the Files Retained counter. Click **Delete Existing Dump Files** if you want to remove existing dump files from your system.

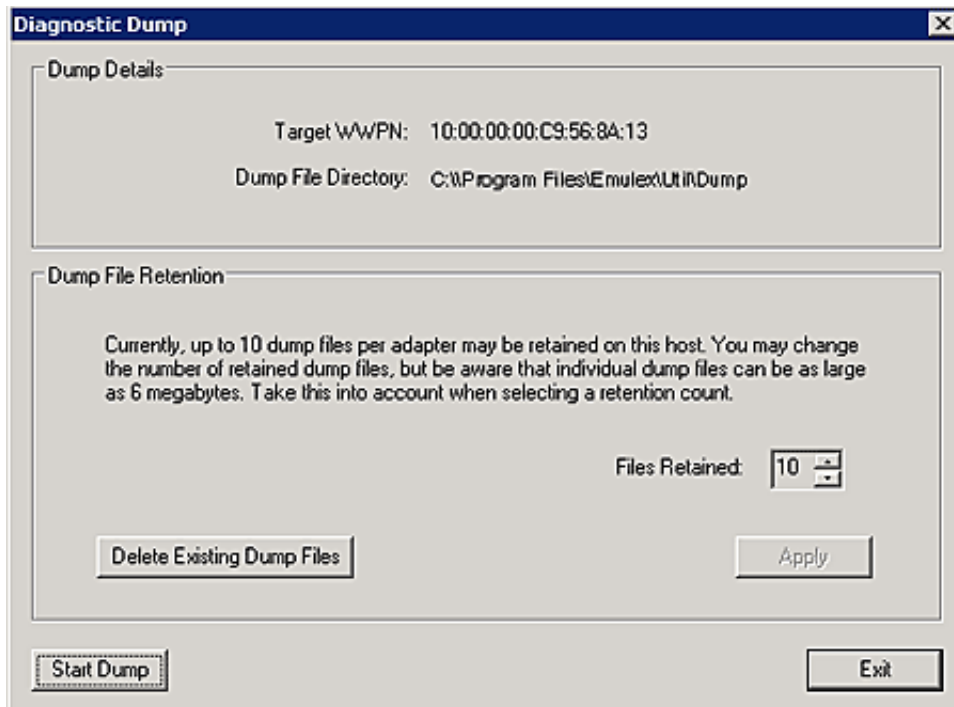


Figure 64: Diagnostic Dump dialog box

3. Click **Start Dump**. Dump files are created. Where these files are created depends upon your operating system:
 - Windows - %ProgramFiles%\Util\Dump\
 - Solaris - /opt/HBAnyware/Dump
 - Linux - /usr/sbin/hbanyware/Dump
 - VMware ESX Server - /usr/sbin/hbanyware/Dump
 Two files are created:
 - <Hostname_WWPN_Date-Time>.DMP
 - <Hostname_WWPN_Date-Time>.TXT

Running Advanced Diagnostic Tests

The Advanced Diagnostics feature gives you greater control than the Quick Test over the type of diagnostics tests that run. Through Advanced Diagnostics, you can specify which tests to run, the number of cycles to run and what to do in the event of a test failure. (Not available in read-only mode.)

Note: Internal and external loopback tests are not available for LP21000 and LP21002 adapters.

To run advanced diagnostics tests:

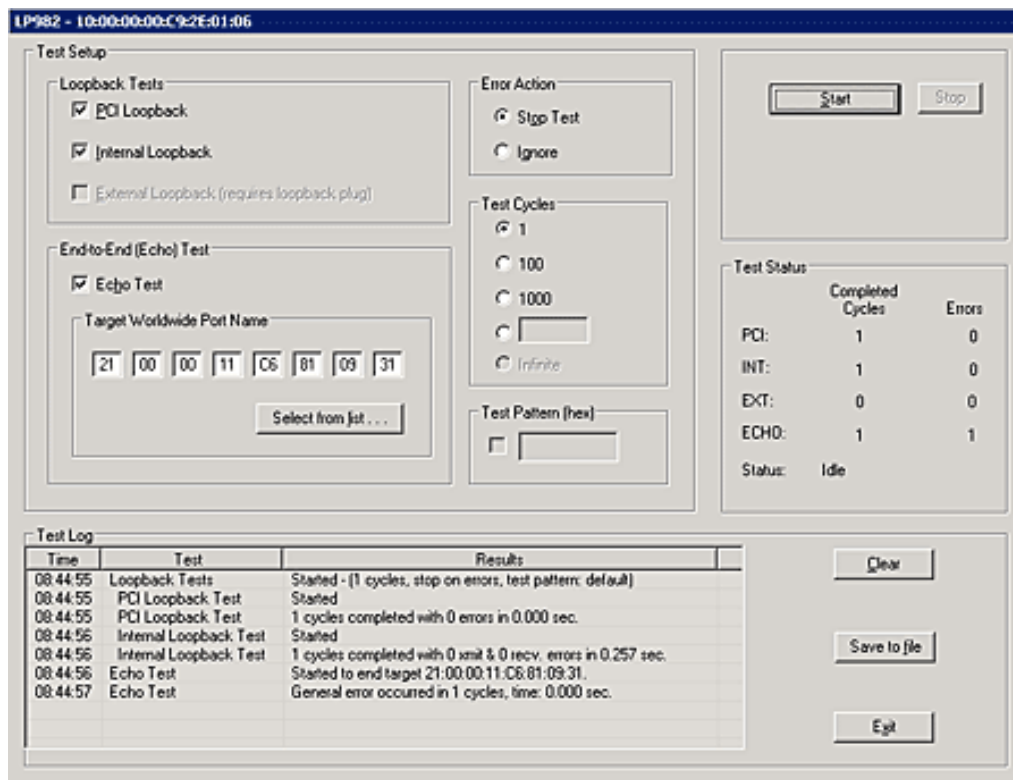
Click **Advanced Diagnostics Test** on the Diagnostics tab to view the Advanced Diagnostics dialog box.

You can run four types of tests:

- PCI Loopback
- Internal Loopback
- External Loopback
- End-to-End (ECHO)

Note: You cannot run the External Loopback test and ECHO test concurrently. If you select External Loopback the ECHO test section is disabled and vice versa.

Test results and the status of running tests are time stamped and appear in the Test Log area.



Test Setup

Loopback Tests

- ☒ PCI Loopback
- ☒ Internal Loopback
- ☐ External Loopback (requires loopback plug)

End-to-End (Echo) Test

- ☒ Echo Test

Target Worldwide Port Name

21 00 00 11 C6 81 09 31

Select from list...

Error Action

- ☒ Stop Test
- ☐ Ignore

Test Cycles

- ☒ 1
- ☐ 100
- ☐ 1000
- ☐ []
- ☐ Infinite

Test Pattern (hex)

☐ []

Test Status

	Completed Cycles	Errors
PCI:	1	0
INT:	1	0
EXT:	0	0
ECHO:	1	1
Status:	Idle	

Test Log

Time	Test	Results
08:44:55	Loopback Tests	Started - (1 cycles, stop on errors, test pattern: default)
08:44:55	PCI Loopback Test	Started
08:44:55	PCI Loopback Test	1 cycles completed with 0 errors in 0.000 sec.
08:44:56	Internal Loopback Test	Started
08:44:56	Internal Loopback Test	1 cycles completed with 0 xmit & 0 rcv, errors in 0.257 sec.
08:44:56	Echo Test	Started to end target 21:00:00:11:C6:81:09:31.
08:44:57	Echo Test	General error occurred in 1 cycles, time: 0.000 sec.

Start Stop

Clear

Save to file

Exit

Figure 65: Advanced Diagnostics

Running Loopback Tests

To run a loopback test, use the Loopback Test section of the Advanced Diagnostics dialog box.

Loopback Test Combinations

Run the following loopback test combinations using the appropriate checkboxes:

- **PCI Loopback Test** - A firmware controlled diagnostic test in which a random data pattern is routed through the PCI Bus without being sent to an adapter link port. The returned data is subsequently validated for integrity.
- **Internal Loopback Test** - A diagnostic test in which a random data pattern is sent down to an adapter link port, then is immediately returned without actually going out on the port. The returned data is subsequently validated for integrity.
- **External Loopback Test** - A diagnostic test in which a random data pattern is sent down to an adapter link port. The data goes out the port and immediately returns via a loopback connector. The returned data is subsequently validated for integrity.

Note: You cannot run the External Loopback test and ECHO test concurrently. If you select External Loopback the ECHO test section is disabled and vice versa.

Error Action

Enables you to define what is to be done in the event of a test failure. There are two error action options:

- **Stop Test** - The error is logged and the test aborted. No further tests run.
- **Ignore** - Log the error and proceed with the next test cycle.

Test Cycles

Enables you to specify test cycles three ways:

- Select an established cycle count by clicking on the corresponding radio button.
- Enter a custom cycle count in the blank field in the Test Cycles area.
- Set the test to run until you manually click Stop, by selecting the Infinite radio button.

Test Pattern

Enter a custom test pattern to be used in tests that transfer data. The test pattern can be up to 8 hexadecimal bytes.

Test Status

The Test Status section displays how many completed cycles of each test ran, as well as the number of errors.

To run loopback tests:

1. From the discovery-tree, select the adapter port on which you want to run the Loopback Test.
2. Select the **Diagnostics** tab and click **Advanced Diagnostics Tests**. From the Loopback Test section of the dialog box, choose the type of Loopback test you want to run and define the loopback test parameters.

Note: You must insert a loopback plug in the selected adapter before running an External Loopback test.

- Click **Start**. The following warning appears:

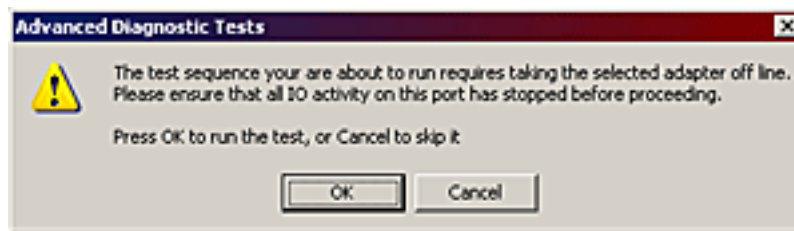


Figure 66: Advanced Diagnostic Tests Warning window

- Click **OK**. If you choose to run an External Loopback test the following window appears:



Figure 67: Advanced Diagnostic Tests Warning window for External Loopback

- Click **OK**. The progress bar indicates that the test is running.

Periodic test feedback, consisting of the current loopback test/cycle plus the completion status of each type of test, is displayed in the Test Log section of the dialog box. Click **Clear** to erase the contents of the log display or click **Save to File** to save the log file.

Running End-to-End (ECHO) Tests

Run echo tests using the End-to-End (ECHO) Test section of the Diagnostics tab. The end-to-end test enables you send an ECHO command/response sequence between an adapter port and a target port. (Not available in read-only mode.)

Note: Not all remote devices respond to an echo command.

You cannot run the ECHO test and the External Loopback test concurrently. If you select the ECHO Test the External Loopback test is disabled.

To run end-to-end echo tests:

- From the discovery-tree, select the adapter port from which to initiate the End-to-End (ECHO) Test.
- Select the **Diagnostics** tab. Click **Advanced Diagnostics Test** (see Figure 68 on page 143).
- Check **Echo Test**. Enter the World Wide Port Name (WWPN) for the target.
or
Click **Select From List** if you do not know the actual WWPN of the test target. The Select Echo Test Target dialog box appears. Select the port to test from the tree-view and click **Select**.

All relevant information for the selected port is automatically added to the Target Identifier section of the Diagnostics dialog box.

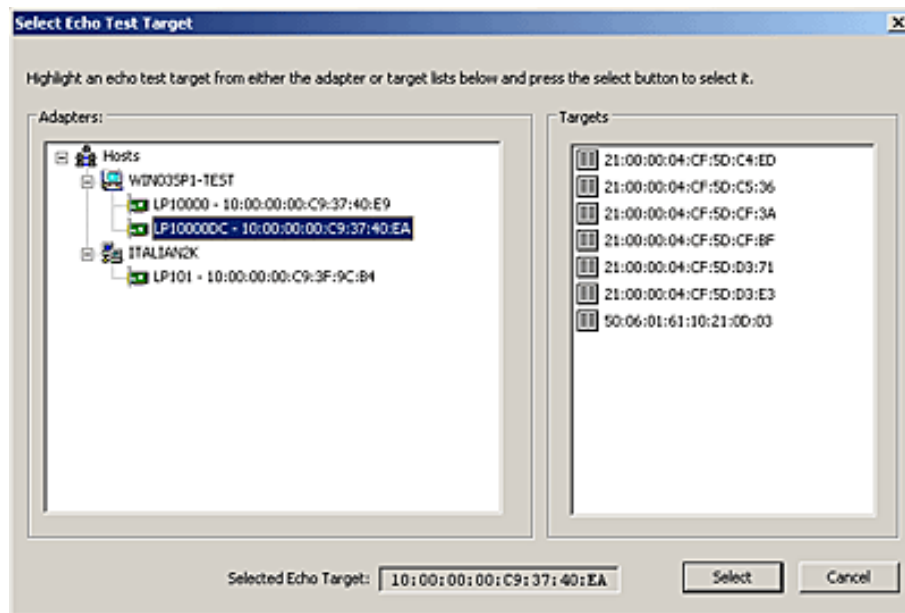


Figure 68: Select Echo Test Target window

4. Define the other parameters you want to use and click **Start**. The following warning window appears:

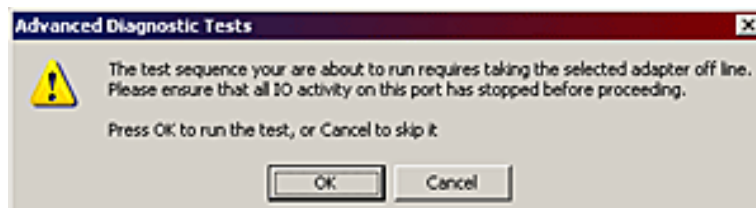


Figure 69: Advanced Diagnostic Tests Warning window

5. Click **OK**. A result screen appears and the test results appear in the Test Log. Click **Clear** to erase the contents of the log display or click **Save to File** to save the log file.

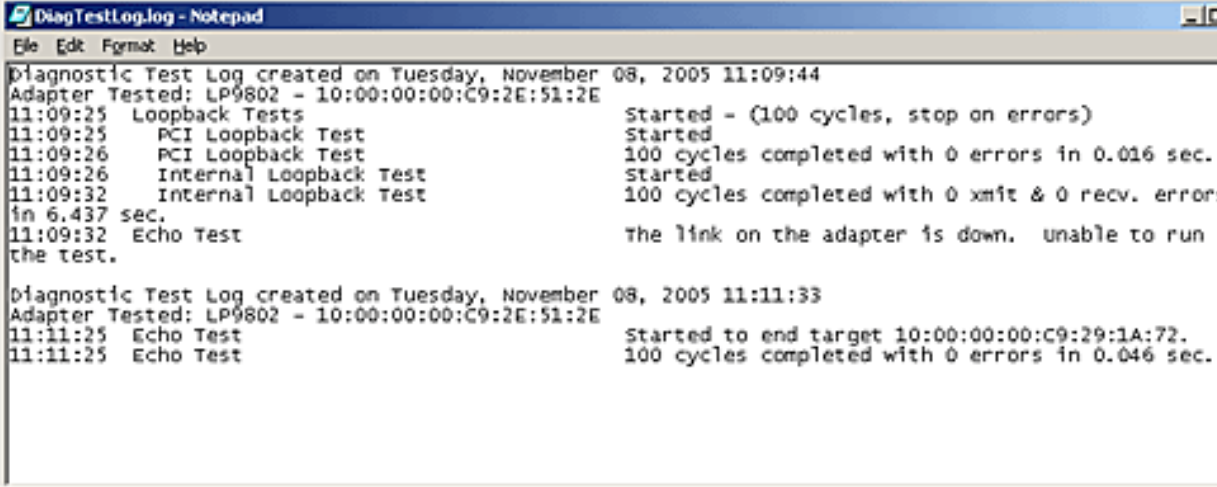
Saving the Log File

You can save the test log to a log file for later viewing or printing. When new data is written to a saved file, the data is appended to the end of the file. Each entry has a two-line header that contains the identifier of the adapter being tested and the date and time of the test. Over time, the data accumulates to form a chronological history of the diagnostics performed on the adapter. (Not available in read-only mode.)

The default location is:

- In Windows: the HBAnyware utility install directory on your local drive
- In Solaris LPFC and Solaris SFS: /opt/HBAnyware/Dump
- In Linux : /usr/sbin/hbanyware/Dump

After writing an entry into the log, you are prompted to clear the display. The default name of the saved file is DiagTestLog.log. An example of a saved log file appears below:



```

DiagTestLog.log - Notepad
File Edit Format Help
Diagnostic Test Log created on Tuesday, November 08, 2005 11:09:44
Adapter Tested: LP9802 - 10:00:00:00:C9:2E:51:2E
11:09:25 Loopback Tests Started - (100 cycles, stop on errors)
11:09:25 PCI Loopback Test Started
11:09:26 PCI Loopback Test 100 cycles completed with 0 errors in 0.016 sec.
11:09:26 Internal Loopback Test Started
11:09:32 Internal Loopback Test 100 cycles completed with 0 xmit & 0 recv. error
in 6.437 sec.
11:09:32 Echo Test the link on the adapter is down. unable to run
the test.

Diagnostic Test Log created on Tuesday, November 08, 2005 11:11:33
Adapter Tested: LP9802 - 10:00:00:00:C9:2E:51:2E
11:11:25 Echo Test Started to end target 10:00:00:00:C9:29:1A:72.
11:11:25 Echo Test 100 cycles completed with 0 errors in 0.046 sec.
  
```

Figure 70: Example of DiagTestLog window

To save the log file:

1. After running a test from the Diagnostic Test Setup dialog box, Click **Save to File**. The Select Diagnostic Log file Name dialog box appears. The default name of a saved file is DiagTestLog.log.
2. Browse to the desired directory, change the log file name if you want and click **Save**.

Changing World Wide Name Configuration

The Maintenance tab enables you to change the World Wide Port Name (WWPN) and the World Wide Node Name (WWNN) of a selected adapter port. For example, you might want to use an installed adapter as a standby in case another installed adapter fails. By changing the standby adapter's WWPN or WWNN it can assume the identity and configuration (e.g. driver parameters, persistent binding settings, etc.) of the failed adapter.


There are three options for referencing WWNs:

- Factory Default WWN - As shipped from the factory.
- Non-Volatile WWN - Values that are saved in non-volatile adapter's flash memory that survives a reboot and/or power outage.
- Volatile WWN - A temporary value that is saved in volatile memory on the flash. If Volatile WWNs are set, they are used instead of the Non-Volatile WWNs.
 - Volatile WWN changes require a warm system reboot in order to take effect. Volatile WWN changes will be lost on systems that power cycle the HBAs during the reboot.
 - Changing volatile WWNs will result in taking the selected adapter offline. Ensure that this adapter is not controlling a boot device and all I/O activity on this adapter has stopped before proceeding. Emulex assumes no responsibility for the consequences of making volatile WWN changes on a boot adapter.

Note: To avoid address conflicts, do not assign a WWNN or WWPN with the HBAnyware utility if you also use another address management tool.

Note: The Change WWN button is disabled for adapters selected on remote hosts running older versions of the HBAnyware utility. The WWPN and WWNN in the Pending Changes area show “n/a” instead of “none”. This also happens when the remote host is busy processing some critical task and WWN Management cannot obtain the current state of WWN management.

To change a port's WWPN or WWNN:

1. Do one of the following:
 - From the **View** menu, click **Hosts**.
 - From the toolbar, click  **Host View**.
2. In the discovery-tree, select the port whose information you want to change.
3. Select the **Maintenance** tab.

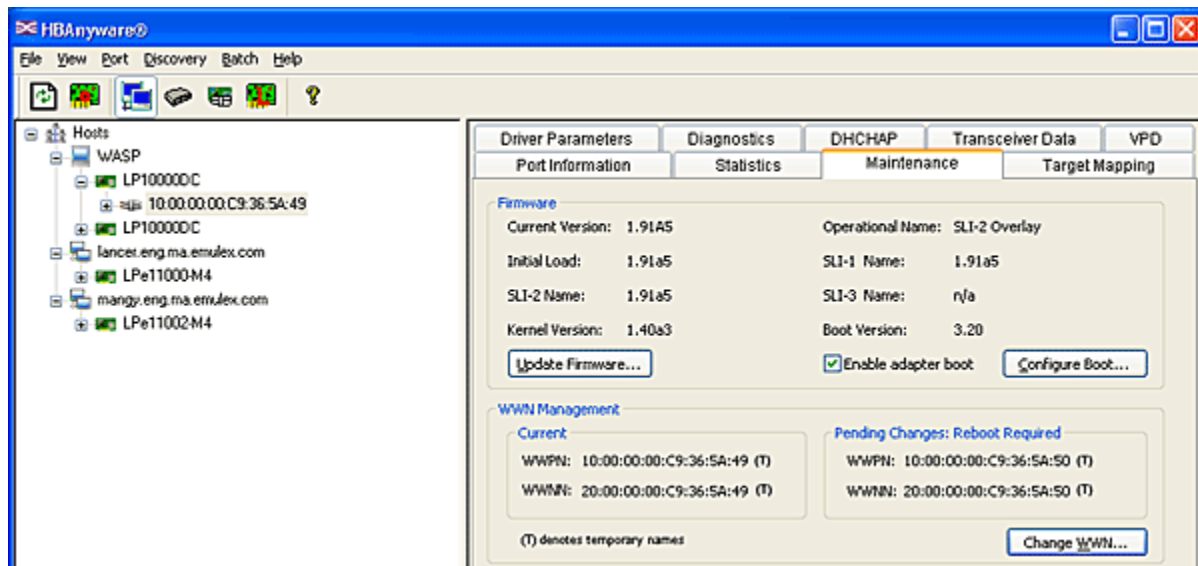


Figure 71: Maintenance tab

4. Click **Change WWN**. The following warning appears:

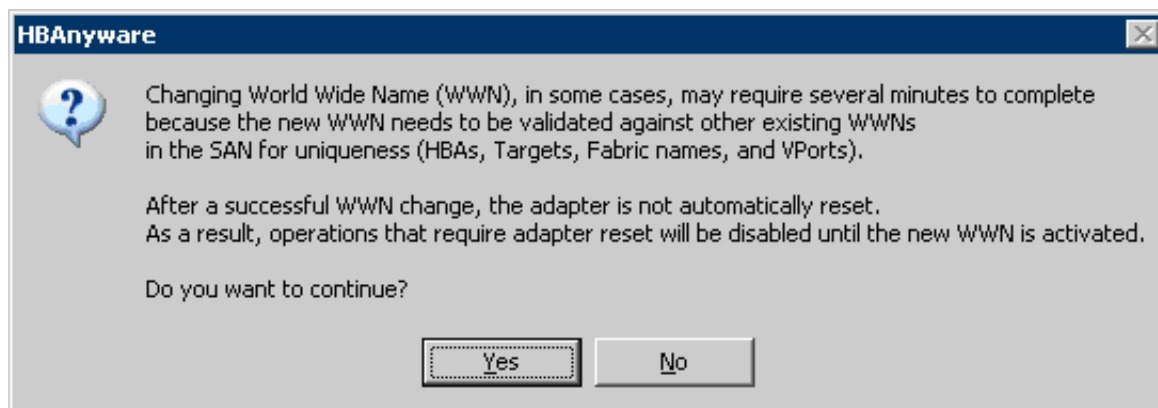
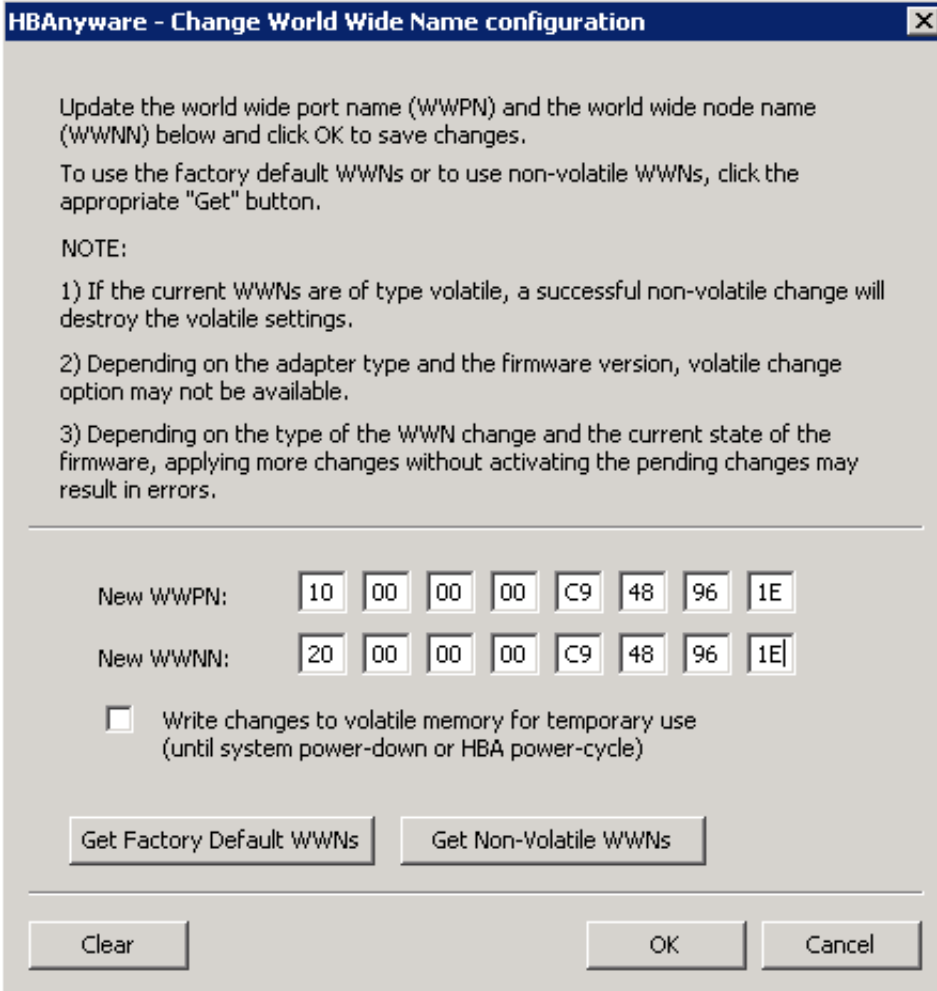


Figure 72: Warning Window About Changing WWN

5. Click **Yes**. The Change World Wide Name configuration dialog box appears.



Update the world wide port name (WWPN) and the world wide node name (WWNN) below and click OK to save changes.

To use the factory default WWNs or to use non-volatile WWNs, click the appropriate "Get" button.

NOTE:

- 1) If the current WWNs are of type volatile, a successful non-volatile change will destroy the volatile settings.
- 2) Depending on the adapter type and the firmware version, volatile change option may not be available.
- 3) Depending on the type of the WWN change and the current state of the firmware, applying more changes without activating the pending changes may result in errors.

New WWPN:

New WWNN:

☐ Write changes to volatile memory for temporary use
(until system power-down or HBA power-cycle)

Figure 73: Change World Wide Name Configuration dialog box

6. Do one of the following:
 - Enter a new WWPN and/or WWNN.
 - Click **Get Factory Default WWNs** to load the settings that were assigned when the adapter was manufactured to the New WWPN and WWNN settings. These values can then be modified if desired and saved as Volatile or Non-Volatile WWNs.
 - Click **Get Non-Volatile WWNs** to load the current Non-Volatile WWN settings to the New WWPN and WWNN settings. These values can then be modified if desired and saved to volatile or non-volatile memory. You can edit the data returned from the button.
7. Check **Write changes to volatile memory for temporary use** to save the New WWPN and New WWNN settings as Volatile WWNs. If unchecked, the New WWPN and New WWNN settings are saved as Non-Volatile WWNs.

Note: If the adapter or firmware does not support Volatile WWNs, the "Write changes to volatile memory for temporary use" check box is disabled. This type of change is supported locally and via TCP/IP connections. This check box is disabled for remote In-band adapters regardless of adapter models and firmware version.

8. Click **OK**. The New WWPN and new WWNN values are saved for Volatile or Non-Volatile use. The new WWPN and WWNN appear in the Pending Changes section in the WWN Management area of the Maintenance tab.
9. Reboot the system for the changes to take effect. The new WWPN and WWNN will appear in the Pending Changes section of the Maintenance dialog box until the system is rebooted. After rebooting, the changes are applied and appear in the Current section of the Maintenance dialog box.

HBAnyware Security

Introduction

After you install the base HBAnyware software, which includes the HBAnyware utility and remote server, on a group of systems, the HBAnyware utility on any of those systems can remotely access and manage the adapters on any systems in the group. This is not a desirable situation because any system can perform actions such as resetting boards or downloading firmware.

You can use the HBAnyware utility security package to control which HBAnyware utility enabled systems can remotely access and manage HBAs on other systems in a FC network. HBAnyware security is systems-based, not user-based. Anyone with access to a system that has been granted HBAnyware client access to remote HBAs can manage those HBAs. Any unsecured system is still remotely accessible by the HBAnyware client software (HBAnyware utility). The HBAnyware security software provides two main security features:

1. Prevent remote adapter management from systems that you do not want to have this capability.
2. Prevent an accidental operation (such as firmware download) on a remote adapter. In this case, you do not want to have access to adapters in systems you are not responsible for maintaining.

When you install the HBAnyware utility security software on a system and run the HBAnyware utility Security Configurator for the first time, that system becomes the Master Security Client (MSC). Only the MSC can view or manage any remote clients. Remote clients can only see the MSC.

Remote clients can manage only by creating an Access Sub-Group (ASG). If you create an ASG, it is then the one and only client, the rest of the machines in the ASG are servers (i.e. servers can not see anybody, only client).

For more information, see “Adding a Server to an ASG” on page 157.

Any system that is already part of the security installation might not run with the proper security attributes if updates to the security configuration are made while it is offline. Any system that is part of the security installation and that is offline when the HBAnyware Security Configurator starts will not be available for security configuration changes even if it is brought online while the Configurator is running.

Starting the HBAnyware Security Configurator

Before starting the HBAnyware Security Configurator:

- Ensure that all of the systems that are part of, or will be part of, the security configuration are online on the Fibre Channel network so that they receive updates or changes made to the security configuration.
- Before running the Security Configurator out-of-band, you must setup the OOB hosts or they will not be seen by the Security Configurator. See the Out-of-Band SAN Management topics for information.
- If you are running the HBAnyware Security Configurator with TCP/IP access, TCP/IP hosts must be added to the discovery-tree or they will not be seen by the Security Configurator.

To start the HBAnyware Security Configurator:

In Windows: On the desktop, click **Start>All Programs>Emulex>HBAnyware Security Configurator**. The HBAnyware Security Configurator Discovery window appears. After discovery is completed, the HBAnyware Security Configurator appears.

To start the HBAware Security Configurator for Linux:

- Run the `/usr/sbin/hbanyware/ssc` script. Type:
`/usr/sbin/hbanyware/ssc`

To start the HBAware Security Configurator for Solaris:

- Run the `/opt/HBAware/ssc` script. Type:
`/opt/HBAware/ssc`

Running the Configurator for the First Time/Creating the ACG

When the HBAware Security software is installed on a system and the HBAware Security Configurator is run for the first time, that system becomes the Master Security Client (MSC). All of the available servers are discovered and available to become part of the system Access Control Group (ACG). You select the systems to be added to the ACG, and the security configuration is updated on all of the selected servers as well as on the initial system. This selection constitutes the participating platforms in this security installation.

To create the ACG:

1. Start the HBAware Security Configurator for the first time in an unsecure environment. The computer from which you run the Configurator will become the MSC. A message is displayed:
2. Click **OK**. The Access Control Group tab is displayed.

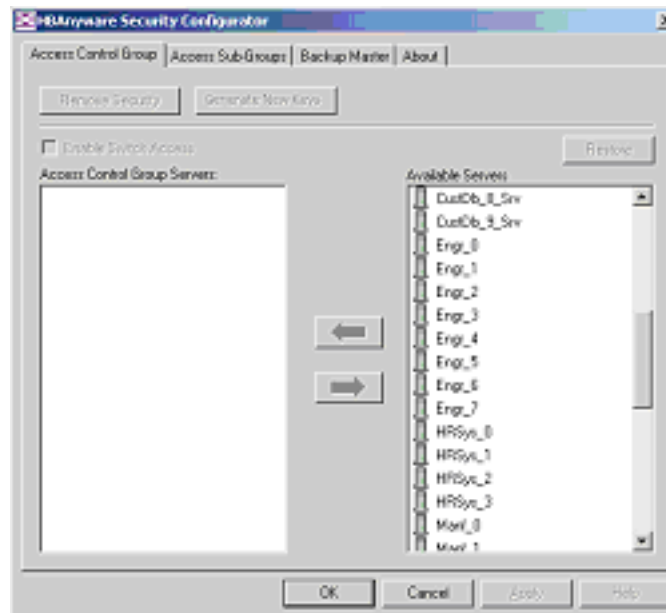


Figure 74: Access Control Group tab - No ACG Servers

3. Select the unsecured servers that you want to add to the ACG from the Available Servers list.

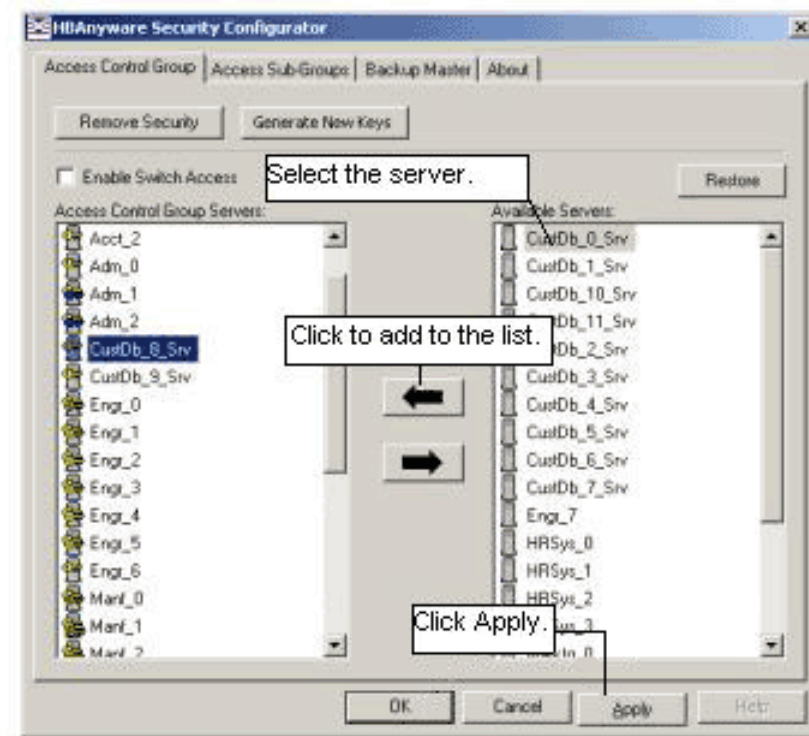


Figure 75: Access Control Group tab with ACG Servers

4. Click the **left arrow** to add the servers to the Access Control Group Servers list.
5. Click **Apply**.

Designating a Master Security Client

The first time you run the HBAware Security Configurator on any system in a FC network, that system becomes the MSC (Master Security Client). See “Running the Configurator for the First Time” on page 150 for more information.

Access Control Groups

Introduction

The Access Control Group tab shows the systems that are part of a client's Access Control Group (ACG) and, from the Master Security Client (MSC), allows you to select the systems that belong to the ACG.

Access Control Group Tab on the MSC

After you have configured security from the MSC for the first time, the Access Control Group tab looks similar to the following:

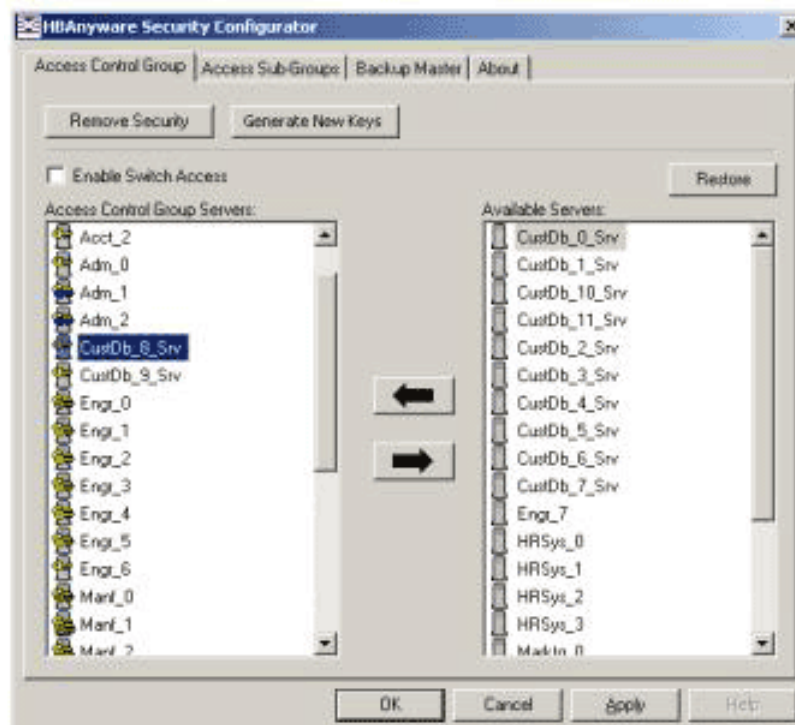


Figure 76: Access Control Group tab on an MSC System

Access Control Group Tab on a Non-MSC

On a non-MSC system, the Access Control Group tab shows the systems that are part of the client's ACG. You cannot modify the ACG on a non-MSC. (You can modify the ACG only on the MSC or a client higher in the security topology's hierarchy.) The ACG tab on a non-MSC system looks similar to the following:

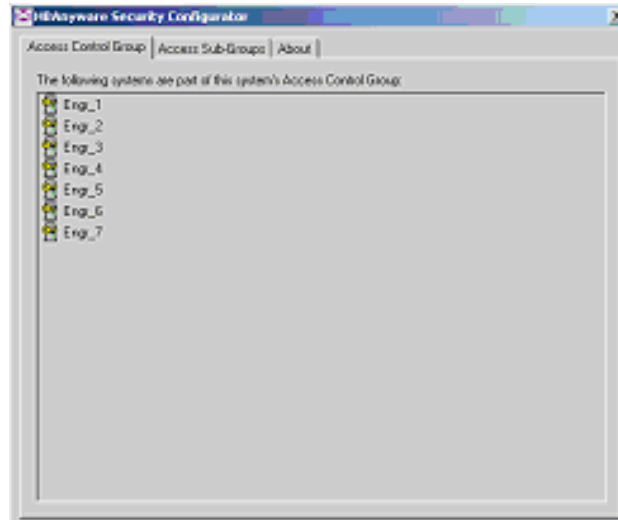


Figure 77: Access Control Group tab on a Non-MSC System

ACG Icons

Depending on the configured security topology, a system can be a server in one or more ACGs. It can also be a client to an ACG. The following icons indicate the state of each of the systems in the Access Control Group Servers list.



The system is a secure server in the ACG. It does not belong to an Access Sub-Group (ASG). You can remove this system from the ACG.



The system is a secure server in the ACG and belongs to one or more ASGs. You can remove this system from the ACG.



The system is a secure server in the ACG and a client to an ASG. You cannot remove this system from the ACG until you remove it as a client from the ASG.



The system is a secure server in the ACG, a secure server in one or more ASGs and a client to an ASG. You cannot remove this system from the ACG until you remove it as a client from the ASGs.



The system is a Backup Master. You cannot remove this system from the ACG until you remove it as a Backup Master.

Adding a Server to the ACG

After you create the initial Access Control Group (ACG) on the Master Security Client (MSC), you can add unsecured servers to the ACG.

To add servers to the ACG:

1. On the Access Control Group tab, from the Available Servers list, select the unsecured servers to add to the ACG (Figure 76).
2. Click the **left arrow** to add the server to the Access Control Group Servers list.
3. Click **Apply**.

Deleting a Server from the ACG

To delete a server from the Access Control Group (ACG):

1. On the Access Control Group tab, from the Access Control Group Servers list, select the secured systems to delete from the ACG (Figure 76).
2. Click the **right arrow** to remove the servers from the Access Control Group Servers list.
3. Click **Apply**.

Removing Security from all Servers in the ACG

You can remove security from all systems only from the Master Security Client (MSC). Removing the entire security topology on all of the servers in the MSC's ACG puts the servers in an unsecure state. The MSC is also put in an unsecure state; consequently, it is no longer the MSC. Any participating systems that are not online will not receive the 'remove security' configuration update, and as a result will no longer be accessible remotely.

To remove security from all servers in the ACG:

1. On the Access Control Group tab, click **Remove Security**. A warning message appears.
2. Click **Yes**. Security is removed from all servers in the ACG.

Generating New Security Keys

You can generate new security keys only from a Master Security Client (MSC). After the new security keys are generated, they are automatically sent to all of the remote servers in the Access Control Group (ACG).

Note: All the servers that are part of the ACG must be online when this procedure is performed so that they can receive the new keys. Any servers that do not receive the new keys will no longer be accessible remotely.

To generate new security keys for all servers in the ACG:

1. From the MSC, start the HBAware Security Configurator. The Access Control Group tab appears (see Figure 74 on page 149).
2. On the Access Control Group tab, click **Generate New Keys**. A dialog box warns you that you are about to generate new security keys for all systems.
3. Click **Yes**. The new keys generate and are sent to all of the remote servers in the ACG.

Restoring the ACG to Its Last Saved Configuration

You can restore the ACG to its last saved configuration, if there are unsaved changes to the ACG, only from the Master Security Client (MSC).

To restore the ACG to its last saved configuration:

From the Access Control Group tab on the MSC, click **Restore** (Figure 76).

Accessing a Switch

You can enable switch access only on a Master Security Client (MSC). Switch access grants the client access rights to a switch to remotely access HBAs on servers in the Access Control Group (ACG).

To enable switch access:

From the Access Control Group tab, check **Enable Switch Access**. (Figure 76).

Access Sub-Groups

Introduction

Use the Access Sub-Group tab to create multiple Access Sub-Groups (ASGs) and multiple levels (tiers) in the security topology hierarchy. The hierarchy can be as many levels deep as desired. However, we recommend the hierarchy extend no more than three levels deep, as it becomes increasingly difficult to keep track of the topology the deeper it goes. The hierarchy shows in the Access Sub-Groups tab as a tree. You can create, modify and delete ASGs at each level in this tree.

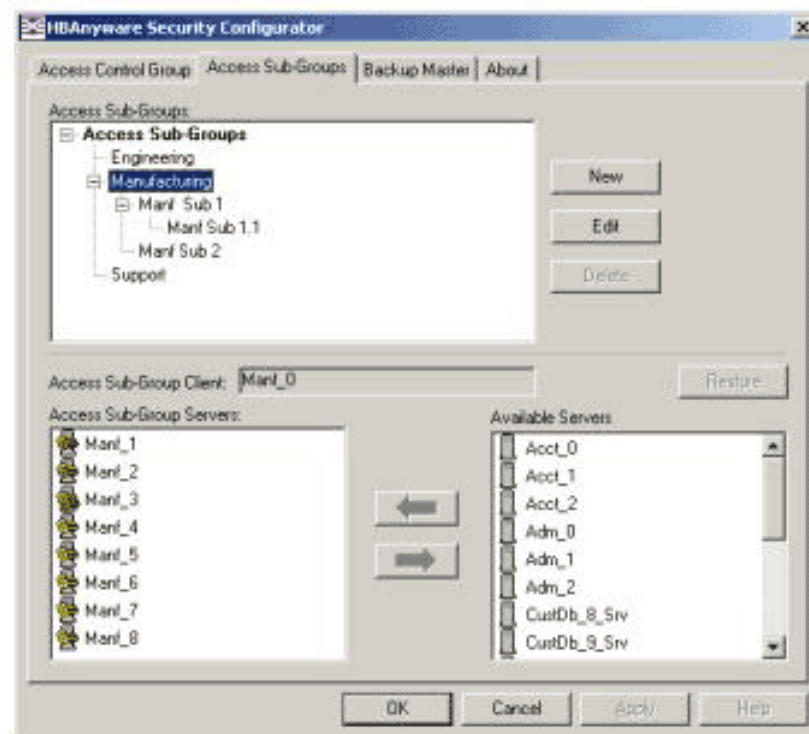


Figure 78: Access Sub-Groups tab with Sub-Groups Created

ASG Icons

The following icons indicate the state of each of the servers in the Access Sub-Group Servers list.



The system is a server in the ASG but not in any child ASGs. You can remove it from the ASG.



The system is a server in the ASG and at least one child ASG. You cannot remove it from the ASG until you remove it from the child ASGs.



The system is a server in the ASG and a client to a child ASG. You cannot remove it from the ASG until you remove it as a client from the child ASG (by either deleting or editing the child ASG).



The system is a server in the ASG, a server in at least one other child ASG and a client to a child ASG. You cannot remove it from the ASG until you remove it from the child ASGs and as a client from the child ASG (by either deleting or editing the child ASG).



The system is a server in the ASG and a client to a non-child ASG. You can remove it from the ASG.



The system is a server in the ASG, a server in at least one child ASG, and a client to a non-child ASG. You cannot remove it from the ASG until you remove it from the child ASGs.

Creating an ASG

After first application of security, nobody can see (remote manage) anybody except for master. Clients are then given ability to remote manage only by ASG creation. What is important but not mentioned here is that, whenever you create any ASG, there is one and only one client, the rest of the machines in the ASG are servers (i.e. servers can not “see” anybody, only client).

Create a new Access Sub-Group (ASG) by selecting one system from the Access Control Group (ACG) to be the client, and some or all of the other systems to be servers to this client, thus defining the new client's ACG. When the HBAnyware Security Configurator is run on the new client, the ACG shows the servers that were configured in the ASG by its parent client.

Note: After first application of security, nobody can see (remote manage) anybody except for master. Clients are then given ability to remote manage only by ASG creation. What is important but not mentioned here is that, whenever you create any ASG, there is one and only one client, the rest of the machines in the ASG are servers (i.e. servers can not “see” anybody, only client)

To create an ASG:

1. Click the **Access Sub-Groups** tab.

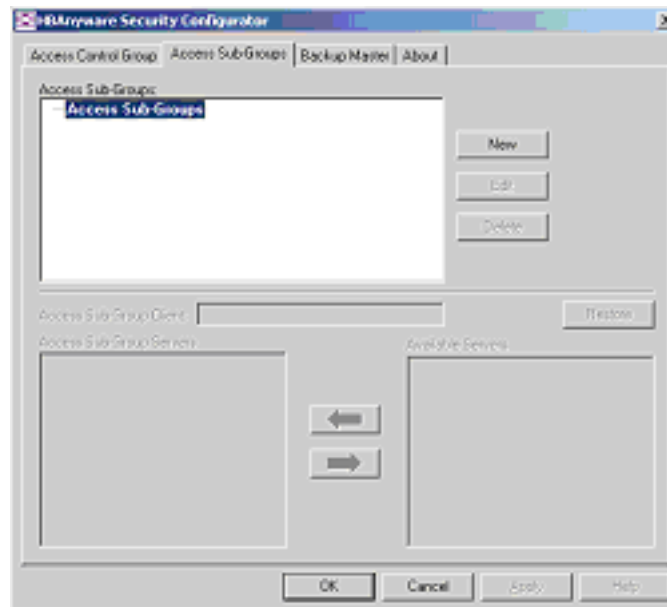


Figure 79: Access Sub-Groups tab with No Sub-Groups Created

2. Click **New**. The New Access Sub-Group dialog box appears:

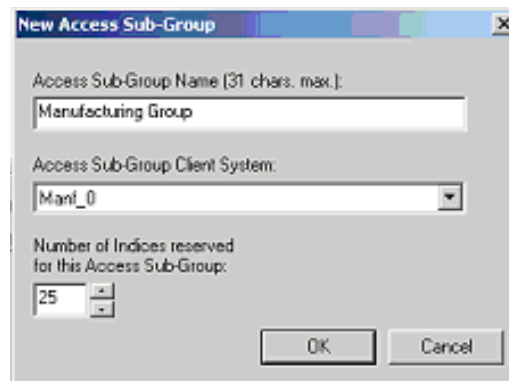


Figure 80: New Access Sub-Group dialog box

3. Enter the ASG information:
 - **Access Sub-Group Name:** Enter the name of the ASG. The ASG name is for identification purposes only. It does not provide any security function. Provide a name that makes it easy to remember the systems that are part of the ASG.

The name can contain any alphanumeric characters, symbols or spaces (up to 31). At each level of the security topology, each ASG name must be unique. If the name is not unique at its level, an error message informs you of this when you click **OK**.
 - **Access Sub-Group Client System:** Select the system that is to be the client.
 - **Number of indices reserved for this Access Sub-Group:** Select the number of 'indices' you want to reserve for the client system of the new ASG. This number reflects the number of subsequent 'child' ASGs that you can subsequently create on the new client's system.
4. Click **OK** in the New Access Sub-Group dialog box. The ASG is created.

Reserved Indices - Examples

A particular security installation can support the creation of several hundred access groups (ACGs and ASGs). When you create each new access group, you allocate some number of 'indices' to the client system of the new ASG. This number reflects the number of subsequent 'child' ASGs that you can subsequently create at the new client's system.

- If zero indices are reserved, you cannot create any lower-level ASG under the client of the new ASG. Thus, if you want to implement a multi-tiered security architecture consisting of many ASGs, and you want to create them all from the Master Security Client (MSC), zero indices would be allocated to each of the new ASGs client platforms when they are created.
- If you create an ASG, and you reserve 25 indices for the new ASG client platform, a child ASG created by this platform has a maximum of only 24 indices available to be reserved (one is taken by the creation of the child ASG itself). This continues down the ASG hierarchy as each lower level ASG is created.
- When you create an ASG from the MSC, a maximum of 50 indices (or less if fewer are available) can be reserved. For all other clients, the maximum depends on how many indices were reserved to that client when its ASG was created, and on how many it has subsequently allocated to its ASGs.

Adding a Server to an ASG

To add a server to an ASG:

1. Click the **Access Sub-Group** tab (see Figure 79 on page 156).
2. The name of the ASG appears in the Access Sub-Groups tree. From the Available Servers list, select the servers to add to the ASG.

Note: TCP/IP accessed servers appear in the Available Servers list even though the ASG client system may not have discovered them yet. These servers can still be added to the Access Sub-Group Servers list.

3. Click the **left arrow** to move the servers to the Access Sub-Group Servers list.
4. Click **Apply** to update servers, adding them to the ASG. The new client can remotely manage the HBAs on those servers using the HBAnyware utility.

Deleting an ASG

Only a leaf node ASG can be deleted (i.e. not ASGs underneath it in the tree). If an ASG has at least one child ASG, you must delete those child ASGs first.

To delete an ASG:

1. From the Access Sub-Group tree, select the leaf node ASG you want to delete.
2. Click **Delete**. A dialog box appears warning you that if you continue the access sub-group will be deleted.
3. Click **Yes**. This operation is immediate. There is no need to click **Apply**.

Restoring an ASG to Its Last Saved Configuration

You can restore an Access Sub-Group (ASG) to its last saved configuration if there are unsaved changes to it.

To restore an ASG to its last saved configuration:

1. Click the **Access Sub-Group** tab (see Figure 79 on page 156).
2. Select the ASG whose configuration you want to restore.

3. Click **Restore**.
4. Click **Apply** to save your changes.

Editing an ASG

You can change the name, client system or reserved indices of an Access Sub-Group (ASG).

To edit an ASG:

1. Click the **Access Sub-Group** tab (see Figure 79 on page 156).
2. Select the ASG you want to edit.
3. Click **Edit**. The Edit Access Sub-Group dialog box appears:

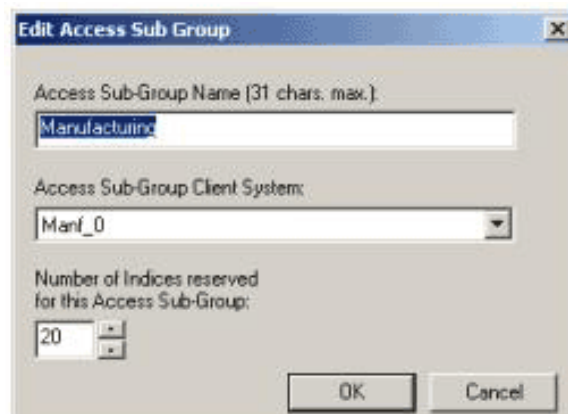


Figure 81: Edit Access Sub Group dialog box

4. Change the ASG information:
 - **Access Sub-Group Name:** Change the name of the ASG. The ASG name is for identification purposes only. It does not provide any security function. Provide a name that logically groups the systems that are part of this ASG.

The name can contain any alphanumeric characters, symbols or spaces (up to 31). At each level of the security topology, each ASG name must be unique. If the name is not unique for its topology level, an error message informs you of this when you click **OK**.

- **Access Sub-Group Client System:** Select the new system to be the client. If the Configurator is running on a system connected to more than one fabric, the client list contains only those systems that can be accessed by the original client of the ASG.
 - **Number of indices reserved for this Access Sub-Group:** Select the new number of 'indices' to reserve for the client system of the new ASG. This number reflects the number of subsequent 'child' ASGs that you can subsequently create on the new client's system. See page 157 for examples.
5. Click **OK** in the Edit Access Sub-Group dialog box to save your changes.

About Offline ASGs

Sometimes a client system is not online when the HBAAnyware Security Configurator is running. In this case, the Access Sub-Group (ASG) for the client appears offline in the ASG tree, much like the following:

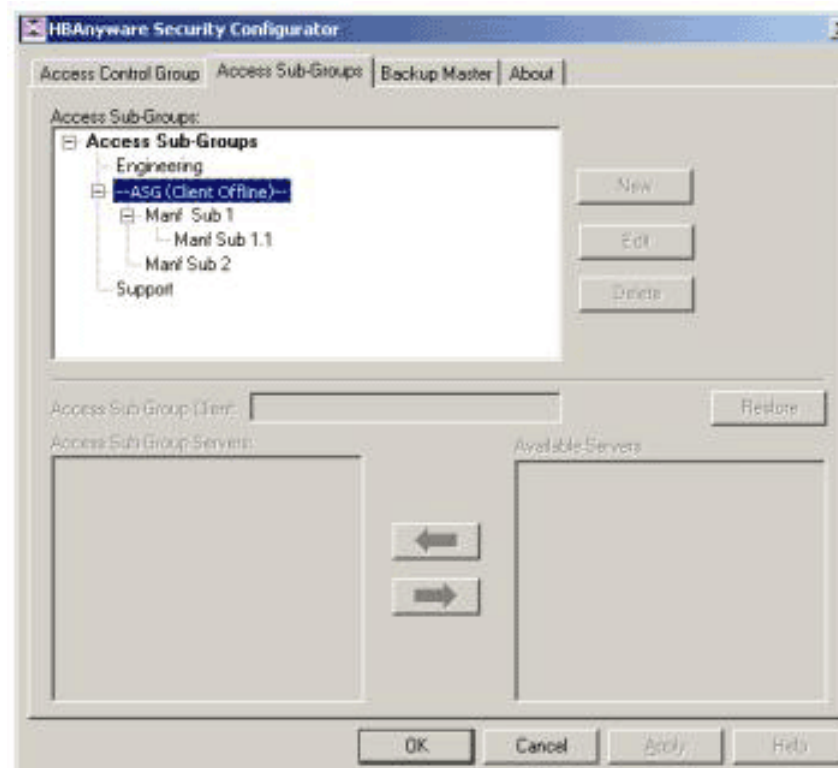


Figure 82: Access Sub-Groups tab - Client System Offline

The offline ASG entry serves as a placeholder for where the real ASG would be in the tree. You cannot modify or delete the entry (although it is removed from the display if all of its child ASGs are deleted).

It is possible to delete the child ASGs of an offline ASG. However, we recommend that you delete them only if the client for the offline ASG is never to be put online again. It is best to delete child ASGs when the parent ASG is online.

If you choose to delete a child ASG, the operation is immediate. There is no need to click **Apply**.

Backup Masters

Introduction

A Backup Master mirrors the security data of the Master Security Client (MSC) in case it has to take over as the MSC if the MSC is unable to operate or is removed from the security configuration. A Backup master system receives all the updates to the security configuration on the MSC. However, you cannot make modifications to the security configuration on a Backup Master.

When the Configurator runs on a Backup Master, the Access Control Group tab looks like the tab on a non-MSC system. The Access Sub-Group tab shows the ASGs, but you cannot change the ASGs (see Figure 76 on page 151).

The Backup Master tab is available only when the HBAware Security Configurator is running on the MSC or a Backup Master. Use this tab to set up a system as a Backup Master to the MSC and to replace the MSC with a Backup Master.

Each time you start the HBAware Security Configurator on the MSC and no Backup Master is assigned, a message warns you that no Backup Master Client is assigned to the security configuration.

If you run the HBAware Security Configurator on a Backup Master, a message warns you that you can only view security information on a Backup Master. Security changes must be made to the MSC.

A Backup Master system receives all the updates that the MSC makes to the security configuration, therefore it is very important that the Backup Master is online when the HBAware Security Configurator is running on the MSC. Otherwise, updates to the security configuration are not reflected to the Backup Master. If the Backup Master then becomes the MSC, the security configuration can be corrupted.

Backup Master Eligible Systems

To be eligible to become a Backup Master, a system must not be a client or server in any ASG. In other words, it must be either a server in the MSC's Access Control Group (ACG) or an unsecure system. If it is an unsecure system, it will be secure when it becomes a Backup Master.

Backup Master Tab and Controls

The first time you select the **Backup Master** tab on the MSC, it looks similar to the following:

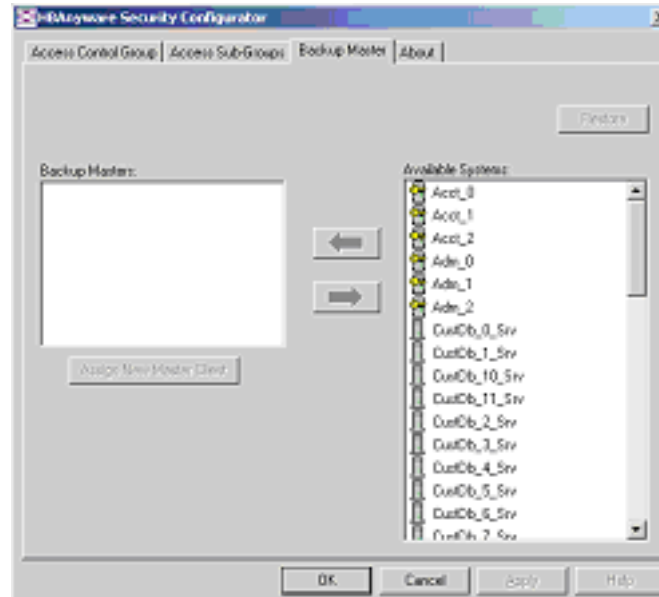


Figure 83: Backup Master tab - First Time Selected

Creating a Backup Master

To create a Backup Master:

1. On the Master Security Client (MSC), start the HBAware Security Configurator.
2. Click the **Backup Master** tab.

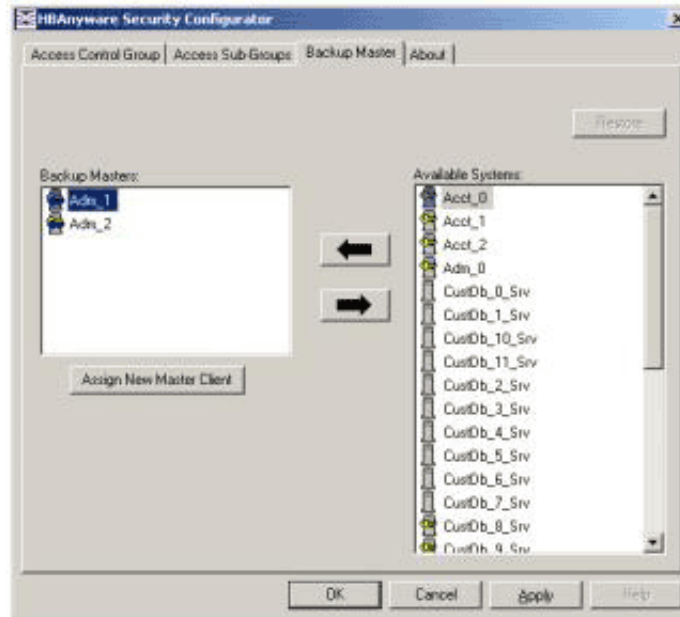


Figure 84: Backup Master tab with Backup Masters

3. Select a system from the Available Systems list.
4. Click the **left arrow** to move the system to the Backup Masters list.
5. Click **Apply** to save your changes.

Reassigning a Backup Master as the New MSC from the Old MSC

Because a Backup Master may have to take over as the Master Security Client (MSC), it must be able to physically access all of the adapters that the MSC can access. If the MSC connects to multiple fabrics, select its Backup Master from the Available Systems list connected to the same fabrics as the MSC.

To reassign a Backup Master as the new MSC from the old MSC:

1. On the current MSC, start the HBAware Security Configurator.
2. Click the **Backup Master** tab (see Figure 84). In the Backup Masters list, select the Backup Master system that you want to reassign as the MSC.
3. Click **Assign New Master Client**. A dialog box appears and asks if you want to proceed.
4. Click **Yes** on the dialog box. The selected Backup Master becomes the new MSC. The current MSC becomes a server in the new MSC's ACG. After the changes are made, a message indicates that the reassignment is complete.
5. Click **OK**. The Configurator closes because the system is no longer the MSC.

Reassigning a Backup Master as the New MSC from the Backup Master

WARNING: Use this method only if the MSC cannot relinquish control to a Backup Master, for example, if you can no longer boot the MSC or connect to the FC network. Under any other circumstances, if the Backup Master takes over as the MSC, and the MSC is still running or comes back online later, there will be two MSCs for the same security configuration. This eventually leads to corruption of the security configuration.

To reassign a Backup Master as the new MSC from the Backup Master:

1. On the Backup Master system that you want to reassign as the MSC, start the HBAAnyware Security Configurator.
2. Click the **Backup Master** tab.

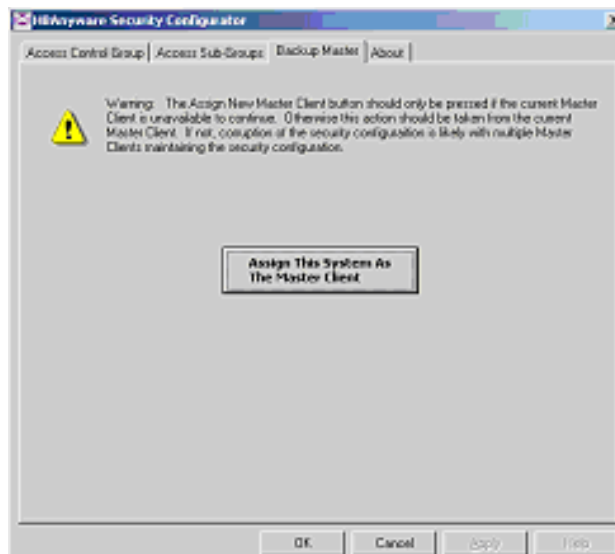


Figure 85: Backup Master “Warning” dialog box

3. Click **Assign This System As The Master Client**. A prompt asks if you want to continue.
4. Click **Yes**. A prompt notifies you that this system is now the new MSC.
5. Click **OK**. The Configurator closes.
6. Restart the HBAAnyware Security Configurator to run the former Backup Master as the MSC.

Using the HBAnyware Utility Command-Line Interface

The Command Line Interface (CLI) Client component of the HBAnyware utility provides access to the capabilities of the Remote Management library from a console command prompt. This component is intended for use in scripted operations from within shell scripts or batch files. The CLI Client is a console application named `hbacmd`. Each time you run this application from the command line, a single operation is performed.

The first parameter of this command is the requested operation. When the specified operation is completed, the command prompt is displayed. Most operations retrieve information about an entity on the SAN and display that information on the console.

Most of the CLI Client commands require one or more additional parameters that specify the nature of the command. A parameter used by many `hbacmd` commands specifies the World Wide Port Name (WWPN) of the adapter that is the target of the command.

For example, run the following command from the directory in which HBAnyware is installed to display the port attributes for the adapter with the specified WWPN:

```
hbacmd portattrib 10:00:00:00:c9:20:20:20
```

`hbacmd` can be run in TCP/IP mode by making the first argument `h=<host>`. For example:

```
hbacmd h=cp-hp5670 listhbas
hbacmd h=138.239.91.121 listhbas
```

Note: For the VMware ESX Server, the firewall on the ESX Server must be opened to manage systems remotely. To enable TCP port #23333, run the following commands:

```
esxcfg-firewall --openPort 23333,tcp,in,hbanyware
esxcfg-firewall --openPort 23333,tcp,out,hbanyware
```

To verify that the correct port is open, run the following command:

```
esxcfg-firewall -q
```

The TCP port number can be changed. If it is not changed, the default is 23333.

Refer to the VMware Server Configuration Guide for more details on how to configure the ESX firewall.

Using the CLI Client

Syntax Rules

The syntax rules for `hbacmd` are as follows:

- All CLI Client commands and their arguments are not case sensitive.
- The requested operation must contain at least three characters, or as many as needed to distinguish it from any other operation.
- Whenever a WWPN is specified, individual fields are separated by colons (:) or spaces (.). When using space separators, the entire WWPN must be enclosed in quotes ("").

The CLI Client Command Reference

CLI Client commands are supported for Windows, Solaris LPFC, Solaris SFS and Linux. Only CLI Client commands that are dynamic are supported for VMware ESX Server.

Note: The PersistentBinding, SetPersistentBinding, RemovePersistentBinding, RemoveAllPersistentBinding, BindingCapabilities, BindingSupport and SetBindingSupport commands are not supported for Linux.

Note: The BindingCapabilities, BindingSupport, GetLunList, PersistentBinding, RescanLuns, RemoveAllPersistentBinding, RemovePersistentBinding, RemoveAllPersistentBinding, SetPersistentBinding, BindingCapabilities, SetBindingSupport, SetLunMask and SetPersistentBinding commands exist in the Emulex driver for ESX Server 3.5.0, but are not supported.

Note: The GetLunMaskbyHBA and GetLunMaskbyTarget commands do not exist for ESX Server 3.5.0.

Read-Only Mode

The CLI (HBACMD) does not allow execution of “sensitive” commands when the HBAnyware utility is configured for read-only mode. An error message will be displayed if such a command is attempted: Error: Read-only management mode is currently set on this host. The requested command is not permitted in this mode.

Help Commands

Help

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: HbaCmd Help

Description: Shows a list of all help commands for the HBAnyware CLI Client application.

Parameters: None

Help Boot

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: HbaCmd Help Boot

Description: Shows a list of all help commands for the boot commands.

Parameters: None

Help BootParams

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd Help BootParams <Parameter Name>

Description: Shows a summary of parameter settings for the adapter and the boot device. Several parameters have detailed help available.

`hbacmd Help BootParams <parameter name>`

Parameter Name (optional) - Specify one of the following boot parameters: AutoScan, BootTargetScan, DevicePathSelection, LinkSpeed, PlogiRetryTimer, or BootParams Topology.

Help GetBootParams

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd Help GetBootParams

Description: Shows help for the GetBootParams command.

Parameters:

WWPN - World Wide Port Name of Object adapter.

Type - None

Help SetBootParams

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd Help SetBootParams

Description: Shows help for the SetBootParams command.

Parameters: None

Attributes Commands

HBAAttributes

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd HBAAttributes <WWPN>

Description: Shows a list of all adapter attributes.

Parameters:

WWPN - World Wide Port Name of the adapter whose attributes you want to view.

PortAttributes

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd PortAttributes <WWPN>

Description: Shows a list of all port attributes for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose port attributes you want to view.

PortStatistics

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd PortStatistics <WWPN>

Description: Shows all port statistics for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose port statistics you want to view.

ServerAttributes

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd ServerAttributes <WWPN>

Description: Shows a list of server attributes for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose server attributes you want to view.

Authentication Commands

DeleteAuthConfig

Supported by: Windows, Solaris LPFC, Solaris SFS and Linux

Syntax: hbacmd DeleteAuthConfig <WWPN1> <WWPN2> <PasswordType> <Password>

Description: Deletes the authentication configuration on the adapter.

Parameters:

WWPN1 - World Wide Port Name of the adapter whose authentication configuration you want to delete.

WWPN2 - Must be ff:ff:ff:ff:ff:ff:ff

PasswordType - 1 = ASCII, 2 = Hex (binary), 3 = Password not yet defined

Password - Current password value.

GetAuthConfig

Supported by: Windows, Solaris LPFC, Solaris SFS and Linux

Syntax: hbacmd GetAuthConfig <WWPN1> <WWPN2>

Description: Retrieves the authentication configuration for the adapter.

Parameters:

WWPN1 - World Wide Port Name of the adapter whose configuration data you want to retrieve.

WWPN2 - Must be ff:ff:ff:ff:ff:ff:ff

InitiateAuth

Supported by: Windows, Solaris LPFC, Solaris SFS and Linux

Syntax: hbacmd InitiateAuth <WWPN1> <WWPN2>

Description: Initiates the authentication configuration on the adapter.

Parameters:

WWPN1 - World Wide Port Name of the adapter whose authentication configuration you want to initiate.

WWPN2 - Must be ff:ff:ff:ff:ff:ff:ff

SetAuthConfig

Supported by: Windows, Solaris LPFC, Solaris SFS and Linux

Syntax: hbacmd SetAuthConfig <WWPN1> <WWPN2> <PasswordType> <Password> <Parameter> <Value>

Description: Sets the authentication configuration for the adapter.

Parameters:

WWPN1 - World Wide Port Name of the adapter whose authentication configuration you want to set.

WWPN2 - Must be ff:ff:ff:ff:ff:ff

PasswordType - 1 = ASCII, 2 = Hex (binary), 3 = Password not yet defined

Password - Current password value

Parameter - Parameters include Mode, Timeout, Bi-directional, Hash-priority, DH-priority, Re-authentication, Re-authentication-interval

Value - Parameter-specific value: Mode = <disabled, enabled, passive>, Timeout = time in seconds, Bi-directional = <disabled, enabled>, Hash-priority = <md5, sha1> (md5 = first md5, then sha1; sha1 = first sha1, then md5), DH-priority = <1,2,3,4,5>, any combination up to 5 digits, Re-authentication = <disabled, enabled>, Re-authentication-interval = < 0, 10 - 3600>

SetPassword

Supported by: Windows, Solaris LPFC, Solaris SFS and Linux

Syntax: hbacmd SetPassword <WWPN1> <WWPN2> <Flag> <Cpt> <Cpw> <Npt> <Npw>

Description: Sets the password for the adapter.

Parameters:

WWPN1 - World Wide Port Name of the adapter for which you want to set a password.

WWPN2 - Must be ff:ff:ff:ff:ff:ff

Flag - 1 = Local (password used by adapter when adapter authenticates to the switch), 2 = Remote (password used by adapter when switch authenticates to the adapter)

Cpt - Current password type is 1 = ASCII or 2 = Hex (binary), 3 = Password not yet defined

Cpw - Current password value.

Npt - New password type is 1 = ASCII or 2 = Hex (binary)

Npw - New password value

Boot Commands

<...> = Required, [...] = Optional

EnableBootCode

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbaCmd EnableBootCode <WWPN> <Flag>

Description: Enables or disables the boot code on the adapter. If the boot code is disabled, the adapter will not boot from SAN, regardless of the value for the EnableBootFromSan boot param. If it is enabled, the adapter will boot from the SAN if the EnableBootFromSan parameter is also enabled.

Parameters:

WWPN - World Wide Port Name of Object adapters

Flag - E = Enable the boot code, D = Disable the boot code

GetBootParams

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbaCmd GetBootParams <WWPN> <Type>

Description: Shows the boot parameters. If any arguments are missing or invalid, a suitable error is reported. If all arguments are ok, the appropriate RM_GetBootParamsXX call is made, and the data is displayed in tabular form.

Parameters:

WWPN - World Wide Port Name of Object adapter.

Type - X86, EFI, OB

SetBootParam

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbaCmd SetBootParam <WWPN> <Type> <Param> <Value1> [BootDev <Value2>]

Description: Performs a high-level read-modify-write operation.

- For Adapter Params, the BootDev keyword and value must be omitted; otherwise, an error is reported.
- For Boot Device Params (OpenBoot) the BootDev keyword and value must be omitted; otherwise, an error is reported.
- For Boot Device Params (X86 and EFI) the BootDev keyword and value are required.

Parameters:

WWPN - World Wide Port Name of Object adapter.

Type - X86, EFI, OB

Param - Parameter Name

Value1 - Parameter Value

Value2 - Boot Device Entry Number: { 0 - 7 }

CEE Commands

Note: CEE commands are for CEE management of LP21000-M and LP21002-M HBAs only.

CEEDownload

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd CEEDownload <WWPN> <Filename>

Description: Updates the CEE firmware on the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter.

Filename - Name of the file to download.

GetCEEPParams

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd GetCEEPParams <WWPN>

Description: Shows the current CEE parameters.

Parameters:

WWPN - World Wide Port Name of the adapter.

SetCEEPParam

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd SetCEEPParam <WWPN> <Param> <Value>

Description: SetCEEPParam configures one of the CEE parameters.

Parameters:

Pausetype - 1 = Standard, 2 = Per Pause Priority

Pfcpriority - 8-bits, each bit representing a pause priority 0 - 7 (e.g. 170 = priorities 7, 5, 3, 1)

Fcoeprriority - 0 - 7

Uifporttype - 1 = Access, 2 = Trunk

Diagnostic Commands

EchoTest

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd EchoTest <WWPN Source> <WWPN Destination> <Count> <StopOnError> <Pattern>

Description: Runs the echo test on adapters.

Note: Support for remote adapter is TCP/IP access only. The EchoTest command fails if the target WWPN does not support the ECHO ELS command.

Parameters:

Source WWPN - World Wide Port Name of the originating adapter.

Destination WWPN - World Wide Port Name of the destination (echoing) adapter.

Count - Number of times to run the test. 0 = run test infinitely

StopOnError - Should the test be halted on Error? 0 = No halt, 1 = Halt

Pattern - Hexadecimal data pattern to transmit (up to 8 characters)

GetBeacon

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd GetBeacon <WWPN>

Description: Shows the current beacon status for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose current beacon you want to view.

LoadList

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd LoadList <WWPN>

Description: Shows the flash load list data for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose flash load list data you want to view.

Loopback

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd loopback <WWPN> <Type> <Count> <StopOnError> <Pattern>

Description: Runs the loop test on the adapter specified by the WWPN.

Note: Only external loopback tests must be run with TCP/IP access.

Note: Internal and external loopback tests are not available for LP21000 and LP21002 adapters.

Parameters:

WWPN - World Wide Port Name of the adapter on which you want to run loopback.

Type - 0 = PCI LoopBack Test, 1 = Internal LoopBack Test, 2 = External LoopBack Test

Count - Number of times to run the test (0 = run test infinitely, Range = 1...99,999)

StopOnError - Should the test be halted on Error? 0 = No halt, 1 = Halt

Pattern - Hexadecimal data pattern to transmit (up to 8 characters).

LoopMap

Supported by: Windows, Solaris LPFC, Solaris SFS and Linux

Syntax: hbacmd LoopMap <WWPN>

Description: Shows the arbitrated loop map data for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose arbitrated loop map data you want to view.

PCIData

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd PCIData <WWPN>

Description: Shows PCI configuration data for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose configuration data you want to view.

HBACMD has a command that displays wakeup parameter information, much the same way that the HBAnyware utility displays it in its own control field.

Wakeup Parameters:

```
Initial Load: 0x02B81991 0x00555637
Flags:        0x00000000
Boot BIOS:    0x03B11713 0x00101303
SLI-1:        0x06B21991 0x00103411
SLI-2:        0x07B21991 0x00103411
Has Expansion Rom: 1
SLI-3:        0x00000000 0x00000000
SLI-4:        0x00000000 0x00000000
Expansion Rom: 0x03B11713 0x00101303
```

The changes suggested for the HBAnyware utility's GUI also apply to this command's output.

PostTest

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd PostTest <WWPN>

Description: Runs the POST on the adapter. Support for a remote adapter is TCP/IP access only.

Parameters:

WWPN - World Wide Port Name of the adapter on which you want to run a POST.

SetBeacon

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd SetBeacon <WWPN> <BeaconState>

Description: Sets the current beacon status for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose beacon you want to change.

BeaconState - New state of the beacon: 0 = Off, 1= On

Wakeup

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd Wakeup <WWPN>

Description: Shows wakeup parameter data for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose wakeup parameter data you want to view.

Driver Parameter Commands

Note: Whenever you chose to set a temporary driver parameter, that is “not permanently”, the parameter is set on each adapter. This method is slightly different then the way it is done for a permanently changed driver parameter. Because of this, the temporarily changed driver parameter must be viewed as an adapter-specific change. To see this change, use `GetDriverParameter` rather than `GetDriverParameterGlobal`. Also, when you run `SaveConfig`, you must run it with the `N` option (adapter-specific). This will gather all the values saved on that HBA. This command must be used cautiously.

DriverConfig

Supported by: Windows, Solaris LPFC, Solaris SFS and VMware ESX Server

Note: For VMware ESX Server: When the `DriverConfig` driver parameter is set persistently and/or requires a reboot, the ramdisk must be rebuilt. To build the ramdisk, type:

```
# esxcfg-boot -b  
# reboot
```

Syntax: `hbacmd DriverConfig <WWPN> <FileName> <Flag>`

Description: Sets all driver parameters for the adapter to the driver parameter values contained in the specified .dpv file type. The .dpv file's driver type must match the driver type of the host platform adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose driver parameters you want to set

FileName - Name of the .dpv file (the file is stored in the Emulex Repository directory)

Flag - G = Make change global (all HBAs on this host), N = Make change non-global (adapter-specific)

GetDriverParams

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server. For VMware ESX Server version 3.5.0 or earlier, the driver uses the `DriverParams` command, but it has the same format as `GetDriverParams`

Syntax: `hbacmd GetDriverParams <WWPN>`

Description: Shows the name and values of each driver parameter for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose driver parameters you want to view.

GetDriverParamsGlobal

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server. For ESX Server version 3.5.0 or earlier, the driver used the `DriverParamsGlobal` command, but it has the same format as `GetDriverParamsGlobal`.

Syntax: `hbacmd GetDriverParamsGlobal <WWPN>`

Description: Shows the name and the global value of each driver parameter for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose driver parameter global names and values you want to view.

SaveConfig

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd SaveConfig <WWPN> <FileName> <Flag>

Description: Saves the specified adapter's driver parameters to a file. The resulting file contains a list of driver parameter definitions in ASCII file format with definitions delimited by a comma. Each definition is of the form: <parameter-name>=<parameter-value>.

Saves either the values of the global set or those specific to the adapter. The file created by this command is stored in the Emulex Repository directory.

Parameters:

WWPN - World Wide Port Name of the adapter whose configuration data you want to save.

FileName - Name of the file that contains the driver parameters list.

Flag - G = Save the global parameter set, N = Save the local (adapter-specific) parameter set

SetDriverParam

Note: For VMware ESX Server: When the DriverConfig driver parameter is set persistently and/or requires a reboot, the ramdisk must be rebuilt. To build the ramdisk, type:
esxcfg-boot -b
reboot

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server.

Description: Allows you to change the value of a driver parameter and designate the scope of that change.

Parameters:

WWPN - World Wide Port Name of the adapter whose driver parameters you want to change.

Flag1 - L = Make change local for this adapter only, G = Make change global (all adapters on this host)

Flag2 - P = Make change permanent (persists across reboot), T = Make change temporary

Note: For VMware ESX Server, CtrlWord - P = Make change permanent, G = Make change global, B = Both, N = Neither. Because P and B are not supported on VMware ESX Server you can only use G or N.

Param - Name of the parameter to modify.

Value - New value you want to assign to the parameter (Input as decimal, prefix with 0x to input as hex).

SetDriverParamDefaults

Note: For VMware ESX Server: When the DriverConfig driver parameter is set persistently and/or requires a reboot, the ramdisk must be rebuilt. To build the ramdisk, type:

```
# esxcfg-boot -b  
# reboot
```

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd SetDriverParamDefaults <WWPN> <Flag1> <Flag2>

Description: Changes all values to the default for the adapter(s).

Parameters:

WWPN - World Wide Port Name of the adapter whose values you want to change to the default.

Flag1 - L = Make changes local for this adapter only, G = Make changes global (all adapters on this host)

Flag2 - P = Make changes permanent (persists across reboot), T = Make changes temporary

Dump Commands

DeleteDumpFiles

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd DeleteDumpFiles <WWPN>

Description: Deletes all diagnostic dump files for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose diagnostic dump files you want to delete.

Dump

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd dump <WWPN>

Description: Displays the maximum number of diagnostic dump files that be can stored for an adapter. Creates a diagnostic dump file in the hbacmd dump file directory.

Parameters:

WWPN - World Wide Port Name of the adapter whose dump information you want to view.

GetDumpDirectory

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd GetDumpDirectory <WWPN>

Description: Displays the dump file directory associated with the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter on which you want to view the dump directory.

GetRetentionCount

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd GetRetentionCount <WWPN>

Description: Displays the maximum number of diagnostic dump files stored for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter on which you want to get the retention count.

SetRetentionCount

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd SetRetentionCount <WWPN> <Value>

Description: Specifies the maximum number of diagnostic dump files stored for the adapter. When the number reaches the retention count limit, the next dump operation causes the oldest diagnostic dump files for that adapter to be deleted.

Parameters:

WWPN - World Wide Port Name of the adapter on which you want to set the retention count.

Value - Value you want to assign to the set retention count.

LUN Masking Commands

Note: The SaveConfig, GetLunList, GetLunMaskbyHBA, GetLunMaskbyTarget, RescanLuns, SetLunMask, DriverConfig, SetDriverParamDefaults and GetAutoConfig commands do not exist for ESX Server.

GetLunList

Supported by: Windows, Solaris LPFC and Solaris SFS

Syntax: hbacmd GetLunList <HBA WWPN> <Target WWPN> <Option>

Description: Queries for the presence of any LUNs.

Parameters:

HBA WWPN - World Wide Port Name of the adapter you want to query.

Target WWPN - World Wide Port Name of the target you want to query.

Option - 0 = Get information from driver, 1 = Get information from configuration

GetLunUnMaskbyHBA

Supported by: Windows, Solaris LPFC and Solaris SFS

Syntax: hbacmd GetLunUnMaskByHBA <HBA WWPN> <Option>

Description: Queries for the presence of any unmasked LUNs by adapter.

Parameters:

HBA WWPN - World Wide Port Name of the adapter you want to query.

Option - 0 = Get information from driver, 1 = Get information from configuration

GetLunUnMaskbyTarget

Supported by: Windows, Solaris LPFC and Solaris SFS

Syntax: hbacmd GetLunUnMaskByTarget <HBA WWPN> <Target WWPN> <Option>

Description: Queries for the presence of any unmasked LUNs by target.

Parameters:

HBA WWPN - World Wide Port Name of the adapter you want to query.

Target WWPN - World Wide Port Name of the target you want to query.

Option - 0 = Get information from driver, 1 = Get information from configuration

RescanLuns

Supported by: Windows, Solaris LPFC and Solaris SFS

Syntax: hbacmd RescanLuns <HBA WWPN> <Target WWPN>

Description: Rescans for the presence of any LUNs.

Parameters:

HBA WWPN - World Wide Port Name of the adapter you want to rescan.

Target WWPN - World Wide Port Name of the target you want to rescan.

SetLunMask

Supported by: Windows, Solaris LPFC and Solaris SFS

Syntax: hbacmd SetLunMask <HBA WWPN> <Target WWPN> <Option> <Lun> <LunCount> <MaskOp>

Description: Masks the specified LUNs.

Parameters:

HBA WWPN - World Wide Port Name of the adapters.

Target WWPN - World Wide Port Name of the target.

Option - 0 = Send information to the driver, 1 = Send information to configuration (make persistent), 2 = Send information to both

Lun - Starting LUN number.

LunCount - Number of LUNs.

MaskOp - A = Mask LUN, B = Clear unmask target level, C = Clear unmask HBA level, D = Unmask LUN, E = Unmask target level, F = Unmask HBA level

Miscellaneous Commands

<...> = Required, [...] = Optional

Download

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd Download <WWPN> <FileName>

Description: Loads the firmware image to the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter to which you want to load firmware.

FileName - File name of the firmware image to load (this can be any file accessible to the CLI client application)

ExportSANInfo

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd exportsaninfo [format]

Note: [format] is optional. Leaving the format parameter blank will store the data in XML format.

Description: For reporting purposes, captures the adapter information in xml or csv format.

Parameters: None

GetVPD

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd GetVPD <WWPN>

Description: Shows the port's Vital Product Data (VPD)

Parameters:

WWPN - World Wide Port Name of the adapter whose VPD you want to view.

ListHBAs

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd ListHBAs

Description: Shows a list of the manageable Emulex adapters discovered by Fibre Channel (in-band) and by TCP/IP (out-of-band).

Parameters: None

Reset

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd Reset <WWPN>

Description: Resets the adapter. An adapter reset can require several seconds to complete, especially for remote devices. Once the reset command is completed, the system command prompt is displayed.

Parameters:

WWPN - World Wide Port Name of the adapter you want to reset.

TargetMapping

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd TargetMapping <WWPN>

Description: Shows a list of mapped targets and the LUNs for the port.

Parameters:

WWPN - World Wide Port Name of the adapter whose target mapping you want to view.

Version

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd Version

Description: Shows the current version of the HBAware CLI Client application.

Parameters: None

Persistent Binding Commands

Note: The PersistentBinding, SetPersistentBinding, RemovePersistentBinding, RemoveAllPersistentBinding, BindingCapabilities, BindingSupport and SetBindingSupport commands are not supported for Linux.

Note: The PersistentBinding, SetPersistentBinding, RemovePersistentBinding, RemoveAllPersistentBinding, BindingCapabilities, BindingSupport and SetBindingSupport commands exist in the Emulex driver for ESX Server, but are not supported.

Note: In order for a binding to take effect immediately (SetPersistentBinding parameter, Scope = I or B), the SCSI Bus and SCSTarget must match the SCSI bus and SCSI target to which the FC target is already automapped. If automapping is disabled, the binding will take effect immediately if the FC target is not already persistently bound and the specified SCSI Bus and SCSTarget are available to be persistently bound. Also, the BindType must match the currently active bind type. Otherwise, you will be notified that you must reboot the system to cause the persistent binding to become active.

AllNodeInfo

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: hbacmd AllNodeInfo <WWPN>

Description: Shows target node information for each target accessible by the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose target node information you want to view.

BindingCapabilities

Supported by: Windows, Solaris LPFC and Solaris SFS

Syntax: hbacmd BindingCapabilities <WWPN>

Description: Shows the binding capabilities present for the adapter. If a binding is configured, it means the binding is maintained across reboots.

Parameters:

WWPN - World Wide Port Name of the adapter whose binding capabilities you want to view.

BindingSupport

Supported by: Windows, Solaris LPFC and Solaris SFS

Syntax: hbacmd BindingSupport <WWPN> <Source>

Description: Shows the binding support available for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose binding support you want to view.

Source - C = Configuration support, L = Live support

PersistentBinding

Supported by: Windows, Solaris LPFC and Solaris SFS

Syntax: hbacmd PersistentBinding <WWPN> <Source>

Description: Specifies which set of persistent binding information is requested: the configured or live state of any present binding.

Parameters:

WWPN - World Wide Port Name of the adapter whose persistent binding information you want to specify.

Source - C = Configuration, L = Live

SetPersistentBinding

Supported by: Windows, Solaris LPFC and Solaris SFS.

Syntax: hbacmd SetPersistentBinding <WWPN> <Scope> <BindType> <TargetId> <SCSIbus>
<SCSITarget>

Description: Sets a persistent binding between an FC target and a SCSI Bus and target. The binding can be to a target WWPN, target WWNN, or target D_ID.

Parameters:

WWPN - World Wide Port Name of the adapter whose persistent bindings you want to set.

Scope - P = Binding is permanent (survives across reboot), I = Binding is immediate, B = Binding is both permanent and immediate.

BindType - P = Enable binding by WWPN, N = Enable binding by WWNN, D = Enable binding by D_ID

TargetId - Target WWPN if BindType = P, Target WWNN if BindType = N, Target D_ID if BindType = D

SCSIbus - Bus number of SCSI device.

SCSITarget - Target number of SCSI device.

RemoveAllPersistentBinding

Supported by: Windows, Solaris LPFC and Solaris SFS

Syntax: hbacmd RemoveAllPersistentBinding <WWPN>

Description: Removes all persisting bindings for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose persistent bindings you want to remove.

RemovePersistentBinding

Supported by: Windows, Solaris LPFC and Solaris SFS

Syntax: hbacmd RemovePersistentBinding <WWPN> <BindType> <ID> <SCSI Bus> <SCSI Target>

Description: Removes persistent binding between an FC target and a SCSI Bus and target. The binding to be removed can be to a target WWPN, target WWNN, or target D_ID.

Parameters:

WWPN - World Wide Port Name of the adapter whose persistent bindings you want to remove.

BindType - P = Remove binding by WWPN, N = Remove binding by WWNN, D = Remove binding by D_ID

ID - Target WWPN if BindType = P, Target WWNN if BindType = N, Target D_ID if BindType = D

SCSI Bus - Bus number of SCSI device.

SCSI Target - Target number of SCSI device.

SetBindingSupport

Supported by: Windows, Solaris LPFC and Solaris SFS

Syntax: hbacmd SetBindingSupport <WWPN> <BindFlag>

Description: Enables and sets the binding support(s) for the adapter.

Parameters:

WWPN - World Wide Port Name of the adapter whose binding support you want to set and enable.

BindFlag - *D = Binding by D_ID, P = Binding by WWPN, * N = Binding by WWNN, *A = Binding by Automap, DA = Binding by D_ID and Automap, PA = Binding by WWPN and Automap, NA = Binding by WWNN and Automap

* Not available for the Storport Miniport driver.

TCP/IP Management Host File Commands

ListHBAs - See "Miscellaneous Commands" on page 179.

Addhost

Supported by: Windows, Solaris LPFC, Solaris SFS and Linux

Syntax: hbacmd addhost host_address

Description: Adds a host to the hosts file. The host_address can be an IP address or a host name.

Parameters:

host_address - Host to add

Removehost

Supported by: Windows, Solaris LPFC, Solaris SFS and Linux

Syntax: hbacmd removehost host_address

Description: Removes a host from the hosts file. The host_address can be an IP address or a host name.

Parameters:

host_address - Host to remove

VPort Commands

<...> = Required, [...] = Optional

CreateVPort

Supported by: Windows, Solaris LPFC, Solaris SFS and Linux

Syntax: hbacmd CreateVPort <physical WWPN> auto [vname]

or

hbacmd CreateVPort <physical WWPN> <virtual WWPN> <virtual WWNN> [vname]

Description: Creates a virtual port with an automatically generated WWPN or a specified virtual WWPN on the specified physical port. If you specify “auto”, the virtual WWPN will be generated automatically. Otherwise, you must specify the virtual WWPN for this parameter. If creation is successful, the WWPN is displayed as part of the output from the command. The optional [vname] parameter can be specified for the virtual port's name.

Parameters:

Physical WWPN - World Wide Port Name of the object adapter.

Virtual WWPN – The virtual World Wide Port Name.

Auto - The virtual WWPN will be automatically generated for the virtual port.

Vname - The virtual port's name (optional).

or

Physical WWPN - World Wide Port Name of the object adapter.

Virtual WWPN – The virtual World Wide Port Name to create.

Vname - The virtual port's name (optional).

DeleteVPort

Supported by: Windows, Solaris LPFC, Solaris SFS and Linux

Syntax: hbacmd deletevport <physical WWPN> <virtual WWPN>

Description: Deletes the virtual port specified by a physical and virtual WWPN.

Parameters:

Physical WWPN - World Wide Port Name of the adapter from which you want to delete a virtual port.

Virtual WWPN - The WWPN for the virtual port.

ListVPorts

Supported by: Windows, Solaris LPFC, Solaris SFS and Linux

Syntax: hbacmd listvports [physical WWPN]

Description: Lists virtual ports on the specified physical port. Leaving the physical wwpn parameter blank will list all VPorts on all manageable hosts that support VPorts.

Parameters:

Physical WWPN - World Wide Port Name of the adapter on which you want to list virtual ports.

VPortTargets

Supported by: Windows, Solaris LPFC, Solaris SFS Linux and VMware ESX Server

Syntax: `hbacmd vporttargets <physical WWPN> <virtual WWPN>`

Description: Lists targets visible to the specified virtual port.

Parameters:

Physical WWPN - World Wide Port Name of the adapter on the targets are visible.

Virtual WWPN - The WWPN for the virtual port.

WWN Management Commands

Note: WWN Management validates WWNs very carefully to avoid name duplication. Therefore, you may see error and warning messages if a name duplication is detected. It is strongly recommended that the activation requirement be fulfilled after each WWN change or restore. When running with “pending changes”, some diagnostic and maintenance features are not allowed.

Change WWNs

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: `ChangeWWN <WWPN> <New WWPN> <New WWNN> <Type>`

Description; Changes the volatile or non-volatile state of WWNs. If the volatile change is requested on an adapter that does not support Volatile WWNs, an appropriate “not supported” error is displayed.

Note: When a volatile change is supported, a reboot is required to activate the new setting. Volatile names will be active until system power-down or adapter power-cycle.

Note: For VMware ESX Server: After changing the WWN of an adapter, be sure your zoning settings are updated before you reboot your ESX server. If the zoning is not updated before your reboot, this could lead to long boot times.

Parameters:

WWPN - World Wide Port Name of Object adapter.

New WWPN - New World Wide Port Name of Object adapter.

New WWNN - New World Wide Node Name of Object adapter.

Type - 0: Volatile, 1: Non-Volatile

Get Capabilities (GetWWNCap on VMware)

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: `hbacmd getwwncap <WWPN>`

Description: Shows if volatile change is supported for the WWPN.

Note: A reboot is required to activate the new setting.

Parameters:

WWPN - World Wide Port Name of Object adapter.

Read WWNs

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: `hbacmd readWWN <WWPN> <Type>`

Description: Reads different types of WWNs.

Parameters:

WWPN - World Wide Port Name of Object adapter.

Type - 0: Volatile, 1: Non-Volatile, 2: Factory Default, 3: Current, 4: Configured

Restore WWNs

Supported by: Windows, Solaris LPFC, Solaris SFS, Linux and VMware ESX Server

Syntax: `RestoreWWN <WWPN> <Type>`

Description: Quickly changes the WWNs back to the factory default or non-volatile values. This change is non-volatile.

Note: A reboot is required to activate the new setting.

Parameters:

WWPN - World Wide Port Name of Object adapter.

Type: 0: Restore Default WWNs, 1: Restore NVRAM WWNs

Troubleshooting

There are several circumstances in which your system may operate in an unexpected manner. The Troubleshooting section explains many of these circumstances and offers one or more workarounds for each situation.

General Situations

Table 16: General Situations

Situation	Resolution
The FC link fails to come up.	Verify that an 8 Gb/s adapter is not attempting to connect to a 1 Gb/s device. Only 2 Gb/s, 4 Gb/s and 8 Gb/s devices are supported on 8 Gb/s HBAs.
The other utilities install, but the HBAnyware utility does not.	<p>You have attempted to install the utilities before installing the Emulex driver.</p> <p>Perform the installation tasks in the following order:</p> <ol style="list-style-type: none"> 1. Install the Emulex driver (see the Installation section of the driver manual). 2. Install the utilities (see the Installation section of the driver manual).
When attempting to start the HBAnyware utility the Web browser displays “Emulex Corporation HBAnyware Demo of HBAnyware WebStart web n.n.n.n...”	<p>The document caching mechanism sometimes behaves erratically if more than one version of Java Runtime is installed on the browser client. There are two workarounds for this problem:</p> <ul style="list-style-type: none"> • Exit the browser and restart it. the HBAnyware utility with Web Launch starts successfully. • Uninstall all non-essential versions of the Java Runtime. HBAnyware Web Launch Service requires that only a single version of the Java Runtime be installed on the browser client. This single version must be JRE version 1.5 or greater.
Operating Error Occurs When Attempting to Run the HBAnyware Utility. When you attempt to run the utility, an operating system error may occur. The computer may freeze.	Reboot the system.
Cannot See Multiple Zones from the Management Server. Cannot see multiple zones on the same screen of my management server running the HBAnyware utility.	Provide a physical FC connection into each of the zones. For each zone you want to see, connect an HBAnyware utility enabled port into that zone. Use Out-of-Band discovery, Ethernet, to connect to the undiscovered server.

Table 16: General Situations (Continued)

Situation	Resolution
<p>Cannot See Other HBAs or Hosts. Although the HBAnyware utility is installed, only local HBAs are visible. The other HBAs and hosts in the SAN cannot be seen.</p>	<p>The utility uses in-band data communication, meaning that the management server running the utility must have a physical FC connection to the SAN. All the adapters in the SAN will be visible if:</p> <ul style="list-style-type: none"> • The other servers have an FC connection to your zone of the SAN. Check fabric zoning. • For Solaris LPFC: All elxhbamgr processes are running on the remote host. To check, enter <code>ps -ef grep elxhbamgr</code>. • All other HBAs are running the HBAnyware utility and the appropriate driver. • The other HBAs are Emulex adapters. <p>Note: The HBAnyware utility must be running on all remote hosts that are to be discovered and managed. Remote capabilities of the HBAnyware utility are subject to fabric zoning configuration. Remote hosts to be discovered and managed by the HBAnyware utility must be in the same zone.</p>
<p>SAN Management Workstation Does Not Have an FC Connection. The SAN management workstation does not have a physical FC connection into the SAN because the other management tools are all out-of-band. Can the HBAnyware utility be run on this SAN management workstation?</p>	<p>The HBAnyware utility can communicate with remote HBAs using out-of-band access as long as the remote host is running the HBAnyware utility.</p> <p>To solve this problem:</p> <ol style="list-style-type: none"> 1. Start the HBAnyware utility. 2. From the Main menu, select Discovery/Out-of-Band/Add Host. The Add Remote Host dialog box appears. 3. In the Add Remote Host dialog box, enter either the name or the IP-address of the host and click OK. When the selected host is discovered, that host and any HBAs running on it will be displayed in the discovery-tree.
<p>Cannot See New LUNs. Although new LUNs were created on the storage array, they do not appear in the HBAnyware utility.</p>	<p>Refresh the screen.</p>
<p>The HBAnyware Security Configurator software package will not install. An error message states that the latest version of the HBAnyware utility must be installed first.</p>	<p>The system either has no HBAnyware software installed or has an older version of the HBAnyware software installed. In either case, obtain the latest version of the HBAnyware software and follow the installation instructions. Remember to install the HBAnyware software before installing the Security Configurator package.</p>
<p>Cannot access formerly accessible servers via the Security Configurator or the HBAnyware utility.</p>	<p>This is actually a symptom of two different problems.</p> <ul style="list-style-type: none"> • New Keys Were Generated While Servers Were Offline • Security Removed While Servers Were Offline <p>See Table 24 on page 198 for details regarding these problems.</p>

Table 16: General Situations (Continued)

Situation	Resolution
Cannot run the Security Configurator on a system that is configured for only secure access. I cannot run the Security Configurator on a system that is configured for only secure server access (it has no client privileges). The following message is displayed when the Security Configurator starts: "This system is not allowed client access to remote servers. This program will exit."	You cannot run the Security Configurator on a system that is configured for only secure server access. Click OK to close the message and the Configurator stops.
Unwanted remote servers appear in the HBAnyware utility.	<p>To prevent remote servers from appearing on the HBAnyware utility, do one of the following:</p> <ul style="list-style-type: none"> • In Windows, disable the HBAnyware service. • In Unix, disable the rmserver process. <p>Disabling this service or process prevents the local servers from being seen remotely.</p>

Emulex Driver for Windows and the HBAnyware Utility Situations

Table 17: Emulex Driver for Windows and the HBAnyware Utility Situations

Situation	Resolution
lputilnt installs, but the HBAnyware Utility Does Not. When you run setupapps.exe, lputilnt installs but the HBAnyware utility does not. You have attempted to manually install the utilities for the driver before manually installing the driver	<p>Perform the installation tasks in the following order:</p> <ol style="list-style-type: none"> 1. Install the driver (see the Installation section of the Emulex Storport Driver User Manual). 2. Install the utilities (see the Installation section of the Emulex Storport Driver User Manual).

Emulex Driver for Solaris LPFC and the HBAnyware Utility Situations

Table 18: Emulex Driver for Solaris LPFC and the HBAnyware Utility Situations

Situation	Resolution
The HBAnyware Utility Appears on Remote Servers in the SAN.	<p>To prevent the HBAnyware utility from appearing on remote servers in the SAN, disable the elxhbamgr process:</p> <ol style="list-style-type: none"> 1. Navigate to /opt/HBAnyware. 2. Run ./stop_hbanyware to stop both the elxhbamgr and elxdiscovry processes. 3. Run ./start_elxhbamgr and ./start_elxdiscovry to restart both processes. <p>Disabling this service or process prevents the local servers from being seen remotely.</p>

Table 18: Emulex Driver for Solaris LPFC and the HBAnyware Utility Situations (Continued)

Situation	Resolution
Cannot See Other HBAs or Hosts.	<p>The HBAnyware utility uses in-band data communication, meaning that the management server running the HBAnyware utility must have a physical Fibre Channel connection to the SAN. All the adapters in the SAN will be visible if:</p> <ul style="list-style-type: none"> • The other servers have a Fibre Channel connection to your zone of the SAN. Check fabric zoning. • Ensure that elxhbamgr processes are running on the remote host: enter <code>ps -ef grep elxhbamgr</code>. • All other HBAs are running the HBAnyware utility and the appropriate driver. • The other HBAs are Emulex adapters. <p>Note The HBAnyware utility must be running on all remote hosts that are to be discovered and managed. Remote capabilities of the HBAnyware utility are subject to fabric zoning configuration. Remote hosts to be discovered and managed by the HBAnyware utility must be in the same zone.</p>

Emulex Driver for Linux and the HBAnyware Utility Situations

Table 19: Emulex Driver for Linux and the HBAnyware Utility Situations

Situation	Resolution
FC link fails to come up	For LP.21000 adapters, ensure the adapter is not in maintenance mode and that it is not running the manufacturing firmware
The HBAnyware software package will not install. An error message states that: "inserv Service Elxlpfc has to be enabled for service ElxDiscSrvinserv: exiting now/sbin/ inserv failed exit code 1."	Reinstall the driver with the lpfc-install script.
If a SAN configuration has 256 targets mapped by the lpfc driver, any additional added targets do not get a target ID mapping by the driver and cause target discovery to fail. Removing targets or reinitializing the link does not solve the problem.	Unload and reload the driver to reset available target IDs. Ensure that the SAN configuration is correct prior to reloading the driver. This will clear the driver's consistent binding table and free target IDs for new target nodes.
In some cases, after loading an OEM supplied combined firmware/OpenBoot image you will not be able to enable BootBIOS from the lputil Boot BIOS Maintenance menu. If you encounter this problem after loading the OEM combined firmware/OpenBoot image, follow the steps outlined in the resolution.	<ol style="list-style-type: none"> 1. Download the current OpenBoot only image for your adapter from the Emulex web site. 2. Load the current OpenBoot only image following steps listed in Updating BootBIOS section of this manual. 3. Run lputil, return to Boot BIOS Maintenance menu. 4. Enable BootBIOS.

Table 19: Emulex Driver for Linux and the HBAware Utility Situations (Continued)

Situation	Resolution
rmmod fails to unload lpfc driver module due to ERROR: Module lpfc is in use. This message can appear when you attempt to remove the driver and there is a Logical Volume Group dependent on the driver.	Make the Logical Volume Group unavailable. Type: <code>lvchange -a n xxxxxxxx</code> where xxxxxx is the Volume Group Name.
LP1005DC-CM2 reported as the LP1050DC. When running lspci or kudzu utilities, you may see the Emulex FC Host Adapter LP1005DC-CM2 reported as the Emulex FC Host Adapter LP1050DC for the pci_id address f0a5. This is due to a delay in getting the pci_id tables updated in the Red Hat and SuSE distributions.	None at this time
An lspci shows recent Emulex HBAs as “unknown”. This is because of the delay of getting new product ID's into the Red Hat and SuSE development cycle.	<p>The VMPilot™ management application (VMPilot 1.2) is a remote-management utility that enhances SAN support for Microsoft Virtual Server using ANSI standard N-Port ID Virtualization (NPIV). VMPilot allows you to create and manage Virtual Ports (VPorts) that provide a virtualized connection to SAN-attached storage.</p> <p>Note: If you use the VMPilot management application on more than one host in the system, version 1.2 must be installed on every host using it. Version 1.2 is not compatible with any earlier version.</p> <p>Note: The HBAware utility can only discover and manage remote HBAs on hosts that are running the HBAware utility's elxhbmgr.</p> <p>For in-band management, remote capabilities of the HBAware utility are subject to fabric zoning configuration. Remote hosts to be discovered and managed by the HBAware utility must be in the same zone.</p>
Slow targets or extended link faults on the storage side may result in storage being marked off-line by the mid-layer and remaining off-line (not recovered) when the link faults are corrected.	This version of the driver should eliminate this problem. However, if you experience off-line device issues, increase the SCSI command timeout to a value greater than or equal to sixty seconds. Emulex also provides a script which addresses this issue (for 2.6 kernels). To access the lun_change_state.sh script, click http://www.emulex.com/support/linux/index.jsp , then click the link to the appropriate driver, and click the Linux tools link.
Under certain conditions of an I/O load, some targets cannot retire an I/O issued by a Linux initiator within the default timeout of 30 seconds given by the scsi midlayer. If the situation is not corrected, the initiator-to-target condition deteriorates into abort/recovery storms leading to I/O failures in the block layer. These types of failures are preceded by a SCSI IO error of hex 6000000.	Emulex provides a script which addresses this issue. To access the set_target_timeout.sh script, click http://www.emulex.com/support/linux/index.jsp , then click the link to the appropriate driver, and click the Linux tools link.

Table 19: Emulex Driver for Linux and the HBAnyware Utility Situations (Continued)

Situation	Resolution
lpfc driver fails to recognize an adapter and logs “unknown IOCB” messages in the system log during driver load. The adapter is running outdated firmware.	Upgrade adapter firmware to minimum supported revision listed in installation guide (or newer).
Loading lpfc or lpfcdfc driver on SLES 9 reports “unsupported module, tainting kernel” in system log.	This message is logged by the SLES 9 kernel whenever a module which is not shipped with the kernel is loaded. This message can be ignored.
rmmod of lpfc driver hangs and module reference count is 0.	Due to a small race condition in the kernel it is possible for an rmmod command to hang. Issue the <code>rmmod -w</code> command. If this does not help, reboot the computer.
System panics when booted with a failed adapter installed.	Remove the failed adapter and reboot.
lpfc driver unload on SLES 9 causes messages like the following to be logged in the system log: “umount: /dev/disk/by-path/pci-0000:02:04.0-scsi-0:0:1:0: not mounted”	These messages are normal output from the SLES 9 hotplug scripts and can be safely ignored.
rmmod fails to unload driver due to Device or resource busy. This message occurs when you attempt to remove the driver without first stopping the HBAnyware utility, when the HBAnyware utility is installed and running or when FC disks connected to a LightPulse adapter are mounted.	Stop the HBAnyware utility before attempting to unload the driver. The script is located in the <code>/usr/sbin/hbanyware</code> directory. Type: <code>./stop_hbanyware</code> Unmount any disks connected to the adapter. Unload the driver. Type: <code>rmmod lpfcdfc</code> Type: <code>rmmod lpfc</code>
Driver Install Fails. The <code>lpfc-install</code> script fails to install the driver.	The install script may fail for the following reasons: <ul style="list-style-type: none"> • A previous version of the driver is installed. Run the <code>lpfc-install --uninstall</code> script and then try to install the driver. • The current driver is already installed. • The kernel source does not match the standard kernel name or you are running a custom kernel.
<p>“No module lpfc found for kernel” error message. When upgrading the kernel, rpm generates the following error: “No module lpfc found for kernel KERNELVERSION”.</p> <p>A recently upgraded kernel cannot find the ramdisk. After upgrading the kernel, the kernel cannot find the ramdisk which halts or panics the system.</p> <p>The driver is not loaded after a system reboot after upgrading the kernel.</p>	<p>These three situations may be resolved by upgrading the kernel. There are two ways to install the driver into an upgraded kernel. The method you use depends on whether or not you are upgrading the driver.</p> <ul style="list-style-type: none"> • Upgrade the kernel using the same version of the driver. • Upgrade the kernel using a new version of the driver. <p>See the Installation section of the driver manual for these procedures.</p>

Table 19: Emulex Driver for Linux and the HBAnyware Utility Situations (Continued)

Situation	Resolution
Driver uninstall fails. The lpfc-install --uninstall script fails with an error.	Try the following solutions: <ul style="list-style-type: none"> Uninstall the HBAnyware utility and SSC software packages. These can be removed by running the ./uninstall script from the HBAnyware utility installation directory. Unmount all FC disk drives. Unload the lpfc and lpfc driver.
lpfc-install script exit code.	The lpfc-install script contains exit codes that can be useful in diagnosing installation problems. See the lpfc-install script for a complete listing of codes and definitions.
The HBAnyware software package will not install. An error message states that: "inserv Service Elxlpfc has to be enabled for service ElxDiscSrvinserv: exiting now/sbin/ inserv failed exit code 1."	Reinstall the driver with the lpfc-install script.
The Emulex driver for Linux does not load in ramdisk for a custom built kernel.	Custom built kernels are not supported by Emulex. However, the Emulex install script will attempt to install the driver into a ramdisk that follows the naming scheme used by Red Hat or SLES kernels. <ul style="list-style-type: none"> The Red Hat naming scheme for IA64 ramdisk images is: /boot/efi/efi/redhat/initrd-KERNELVERSION.img. The Red Hat naming scheme for ramdisk images on all other architectures is: /boot/initrd-KERNELVERSION.img. SLES names follow a similar scheme for IA64. If a custom built kernel has a ramdisk image that does not follow the appropriate naming scheme, the name of the image can be changed using the following procedure: <ol style="list-style-type: none"> Change the name of the ramdisk image to match either the Red Hat or SLES naming scheme, depending on the distribution being used. Update any file links to the HBAnyware utility ramdisk image. Edit the boot loader configuration file: (i.e., /etc/lilo.conf, /etc/yaboot.conf, /boot/grub/grub.conf, /boot/grub/menu.lst), find any references to the old ramdisk image name, and replace them with the new name. Reboot the system to verify the changes. Install the Emulex lpfc Linux driver kit.
The Linux SCSI subsystem only sees 8 LUNs when more are present.	Some SCSI drivers will not scan past 8 LUNs when the target reports as a SCSI-2 device. Force SCSI Bus scan with /usr/sbin/ lpfc/lun_scan. SuSE supplies /bin/rescan-scsi-bus.sh which can be changed to scan everything.
Cannot See Any HBAs. You launch the HBAnyware utility and no adapters are visible.	Try the following solutions: <ol style="list-style-type: none"> Perform an 'lsmod' to see if the Emulex drivers are loaded. Look for an error message on the command line stating the lpfc driver is not loaded. If this is the case, do an insmod of the lpfc lpfc driver and re-launch the HBAnyware utility. Exit the HBAnyware utility and run ./stop_hbanyware. Then run ./start_elxhamgr and ./start_elxdiscovery, and re-launch the HBAnyware utility. The HBAs should be visible. If they are not visible reboot your system.

Table 19: Emulex Driver for Linux and the HBAnyware Utility Situations (Continued)

Situation	Resolution
<p>Cannot See Other HBAs or Hosts. Although the HBAnyware utility is installed, only local adapters are visible. The other adapters and hosts in the SAN cannot be seen.</p>	<p>All the adapters in the SAN will be visible if:</p> <ul style="list-style-type: none"> • The other servers have a connection to your zone of the SAN. Check fabric zoning. • The elxhbamgr processes are running on remote hosts (enter <code>ps -ef grep elxhbamgr</code>). • All other HBAs are running the HBAnyware utility and the appropriate driver. • The other HBAs are Emulex adapters. <p>Note: The HBAnyware utility services must be running on all remote hosts that are to be discovered and managed. If the HBAnyware Security Configurator is running, only the master or Access group client can see the servers.</p>
<p>Cannot See New LUNs. Although new LUNs were created on the storage array, they do not appear in the HBAnyware utility.</p>	<p>Try the following:</p> <ol style="list-style-type: none"> 1. Refresh the screen. 2. Exit the HBAnyware utility and restart it. If new LUNs are visible, you are finished. <p>If that doesn't work, try the following:</p> <ol style="list-style-type: none"> 1. Exit the HBAnyware utility. 2. Navigate to <code>/usr/sbin/hbanyware</code>. 3. Run <code>./stop_hbanyware</code> to stop both the elxhbamgr and elxdiscovry processes. 4. Run <code>./start_elxhbamgr</code> and <code>./start_elxdiscovry</code> to restart both processes. 5. Start the HBAnyware utility.
<p>Unwanted Remote Servers Appear in the HBAnyware utility</p>	<p>To prevent unwanted servers from appearing in the HBAnyware utility, do the following:</p> <ol style="list-style-type: none"> 1. Navigate to <code>/usr/sbin/hbanyware</code>. 2. Run <code>./stop_hbanyware</code> to stop both the elxhbamgr and elxdiscovry processes. 3. Run <code>./start_elxhbamgr</code> and <code>./start_elxdiscovry</code> to restart both processes. Disabling this service or process prevents the local servers from being seen remotely.
<p>Cannot access formerly accessible servers via the Security Configurator or the HBAnyware Utility.</p>	<p>This is actually a symptom of two different problems.</p> <ul style="list-style-type: none"> • New Keys Were Generated While Servers Were Offline • Security Removed While Servers Were Offline <p>See Table 24 on page 198 for details regarding these problems.</p>

VPorts and the HBAnyware Utility Situations

Situation	Resolution
VPort Creation Failure	If an error occurs during VPort creation, an error message indicates the failure.
Virtual Ports for Unsupported Adapter or Host	When you select an unsupported adapter port or host that is running an older version of the HBAnyware utility, "Virtual Ports not available on this HBA or Host". appears in the Virtual Port window.
Port Not Ready	<p>The controls in the New Virtual Port box of the Virtual Port window are replaced by a list of reasons why VPorts cannot be created. The reasons can be one or more of the following:</p> <ul style="list-style-type: none"> Driver NPIV parameter is disabled. SLI-3 is not being used by port. Adapter port is out of resources for additional virtual ports. The port is not connected to a fabric. The fabric switch does not support virtual ports. The fabric switch is out of resources for additional virtual ports. The port link state is down.

Security Configurator Situations - Access Control Groups (ACG)

Table 20: Access Control Groups Situations

Situation	Resolution
<p>All servers are not displayed under one of these two circumstances:</p> <ul style="list-style-type: none"> When I run the Security Configurator on the MSC, I do not see all of the systems in available servers or ACG Servers lists. When I run the Security Configurator on a non-MSC, I do not see all of the systems I should see in the ACG Servers list. 	<p>Make sure all of the systems are connected to the FC network and are online when you start the Configurator. Discovery of the systems is done only once, at startup. Unlike the HBAnyware utility, there is no Discovery Refresh button. Therefore, the Security Configurator must be restarted to rediscover new systems.</p>
<p>Cannot add or remove a server. The Security Configurator shows only a list of the systems in this system's ACG. I cannot add or remove systems from the ACG.</p>	<p>This is normal. You can modify the ACG for your system only on the MSC or on a parent client system.</p>
<p>The HBAnyware utility shows non-ACG Servers. The HBAnyware utility shows servers that are part of the ACG and that are not part of the ACG.</p>	<p>The HBAnyware utility discovers unsecured servers as well as servers that are part of its ACG. The servers you see that are not part of the ACG are unsecured. They are discovered by any system running the HBAnyware utility on the same FC fabric.</p>

Security Configuration Situations - Access Sub-Groups (ASG)

Table 21: HBAware Security Configurator - Access Sub-Groups Situations

Situation	Resolution
<p>Cannot add or remove a server.</p>	<p>When all of the systems in an ACG are running on a single fabric, they are all available to be added to any ASG. However, if the client is connected to more than one fabric, it is possible that not all of the servers in the client's ACG are physically accessible by a chosen client for an ASG. In this case, those servers are not available to be added to that ASG.</p> <p>If you add a system to an ASG as a server, and then make the system a client to a child ASG, you cannot remove it from the ACG it belongs to as a server until you delete the ASG to which it is a client.</p> <p>Before you delete a server from an ASG, you must first remove the server from any lower level ASGs to which it belongs.</p>
<p>In the ASG tree of the Access Sub-Groups tab, one or more of the names of the ASGs is displayed as “- ASG (Client Offline) -”.</p>	<p>The client system for the ASG was not discovered when the Configurator was started. This is actually a symptom of two different problems.</p> <ul style="list-style-type: none"> • All Servers Are Not Displayed • New Keys Were Generated While Servers Were Offline <p>See Table 24 on page 198 for details regarding these problems.</p>
<p>Not All Servers are available to an ASG. When you create a new ASG or modify an existing ASG, not all of the servers in the ACG are available to be added to the ASG.</p>	<p>A client system can be connected to more than one fabric. While the system the Security Configurator is running on can access all of the servers in its ACG, the selected client for the ASG might not have access to all of the servers. Only those that can be accessed by the selected server will be available.</p>

HBAnyware Security Configurator Situations - Backup Masters

Table 22: HBAnyware Security Configurator - Backup Masters Situations

Situation	Resolution
Cannot create a backup master.	<p>Select a system (or group of systems) from the MSC to be the Backup Master. The system must be either an unsecured system (which will be secured by being made a Backup Master), or a system that is not part of any ASG (client or server). These systems will mirror the MSC's security configuration.</p> <p>Because the Backup Master may some day take over as the MSC, the Backup Master must be able to physically access all of the systems that the MSC can access. Therefore, if the MSC is connected to multiple fabrics, the Backup Master also must be connected to those same fabrics. When you select a Backup Master, the HBAnyware Security Configurator displays a warning if it detects that the system selected to be a Backup Master is not able to physically access the same systems that the MSC can access.</p>
Cannot modify the Security Configurator.	<p>Select a system (or group of systems) from the MSC to be the Backup Master. The system must be either an unsecured system (which will be secured by being made a Backup Master), or a system that is not part of any ASG (client or server). These systems will mirror the MSC's security configuration.</p> <p>The Backup Master has client access from the HBAnyware utility to all of the servers in the MSC's ACG. However, the Backup Master does not have client access to the MSC and it cannot modify the security configuration (create, modify or delete ASGs).</p>
No Backup Master and the MSC is no longer available. I do not have a Backup Master and the MSC system is no longer available. The servers are still secure. I installed the Security Configurator on another system, but I cannot access those servers to remove the security from them.	<p>The servers are no longer part of a valid security configuration because there is no MSC to provide master control of the configuration. In order to reset the security on the affected servers, you must contact Emulex Technical Support to receive a special application and instructions on the reset procedure. After the servers have been reset, they should be seen by the Security Configurator and the HBAnyware utility. At this point, you can set up security again through another MSC. At this time, also create a Backup Master.</p>
The Backup Master tab is not available.	<p>The Backup Master tab is displayed only when the Security Configurator is running on the MSC or a Backup Master. You use this tab to set up a system or systems to be backups to the MSC and to replace the MSC with a Backup Master.</p> <p>Each time you start the Security Configurator on the MSC and there is no Backup Master assigned, a warning message urges you to assign at least one Backup Master to prevent the loss of security information if the MSC were to become disabled.</p>

Error Message Situations

Table 23: Error Message Situations

Situation	Resolution
Error Message Appears When Creating an ASG. This message appears when you create an ASG: "The Access Sub-Group name already exists. Please use a different name."	You entered a duplicate ASG name in the Access Sub-Group Name field. At each level of the security topology, each ASG name must be unique. Click OK on the message and enter a unique ASG name.
Error Message Appears When Deleting an ASG. This error message appears when you delete an ASG: "The Access Sub-Group parent's ASG is offline. You should delete the ASG when the parent ASG is available. This ASG should only be deleted if the parent ASG will not be available again. Are you sure you want to delete this ASG? "	The offline ASG entry serves as a placeholder for where the real ASG would be in the tree. You can neither modify nor delete it (although it is removed from the display if all of the child ASGs are deleted). It is possible to delete the child ASGs of the offline ASG. However, it is recommended that you delete them only if the client for the offline ASG will never be put online again. It is best to delete child ASGs when the parent ASG is online. Click Yes on the error message to delete the ASG or No to close the message without deleting.
Error Message Appears When Starting the HBAware Security Configurator. This message appears when you start the Security Configurator: "This system is not allowed client access to remote servers. This program will exit."	The system you are running the Security Configurator on is already under the security umbrella as a server to one or more clients. To make this server a client (so that it can successfully run the Security Configurator), click OK to close the message and exit the program, then do the following: <ol style="list-style-type: none"> 1. Run the Security Configurator on the MSC or on any client that has this server in its ASG. 2. Make this server a client to a group of servers.
Error Message States "No Backup Master Client Assigned". This message appears when you start the Security Configurator: "There are no Backup Master Client Systems assigned to this security configuration. At least one should be assigned to avoid loss of the security configuration should the Master Client System become disabled."	Use the Backup Master tab to assign a Backup Master for the MSC.
Error Message States "Utility is Running on an Unsecure System". This message appears the first time you start the Security Configurator in an unsecure environment: "This utility is running on an unsecure system. Continuing will allow you to set up a new security configuration making this system the Master Client System."	Click OK on the message and complete the ACG setup. The system on which the Security Configurator is running will become the MSC.
Error Message States "System is a Backup Master Client System". This warning appears when you start the Security Configurator on a Backup Master system. "Warning: This system is a backup master client system. Therefore you will only be able to view the security configuration. To make changes, you will need to run this utility on the master client system."	Because each Backup Master system receives all the updates that the MSC makes to the security configuration, the Backup Master systems must be online when the Security Configurator is running on the MSC. Otherwise, updates to the security configuration are not reflected to the Backup Master. If the Backup Master becomes the MSC, corruption of the security configuration may occur. Click OK to close the message.

Master Security Client Situations

Table 24: Master Security Client Situations

Situation	Resolution
<p>The MSC is no longer bootable or able to connect to the FC network.</p>	<p>You must reassign a Backup Master as the new MSC from the Backup Master.</p> <p>Warning: Use this procedure only if the MSC cannot relinquish control to a Backup Master. For example, if the MSC is no longer bootable or able to connect to the FC network. Under any other circumstances, if the Backup Master takes over as the MSC and the MSC is still running or comes back online later, there will be two MSCs for the same security configuration. This will eventually lead to corruption of the security configuration.</p>
<p>New Keys Were Generated While Servers Were Offline. A “Generate New Keys” operation was performed while one or more of the servers were offline. Now those servers can no longer access the HBAnyware Security Configurator or the HBAnyware utility.</p>	<p>The servers are no longer part of the security configuration. In order to reset the security on the affected servers, you must contact Emulex Technical Support to receive a special application and instructions on the reset procedure. After the servers have been reset, they can be added back into the security topology by the MSC.</p> <p>Note: If the server was also a client to an ASG, then when you run the Security Configurator on the MSC or a parent client of this client, its label in the ASG tree of the Access Sub-Group tab will be “- ASG (Offline Client) -”. You must delete the ASG (after deleting the child ASGs) and recreate the ASG configuration of this client and its child ASGs.</p>
<p>Security Removed While Servers Were Offline. Security was removed while one or more servers were offline. I can no longer access those servers from the Security Configurator or the HBAnyware utility.</p>	<p>The servers are no longer part of the security configuration. In order to reset the security on the affected servers, contact Emulex Technical Support to receive a special application and instructions on the reset procedure. After the servers have been reset, they should be seen by the Security Configurator or the HBAnyware utility.</p>